FOURTH EDITION

mixing
audio

CONCEPTS | PRACTICES | TOOLS

ROEY IZHAKI
www.MixingAudio.com

ROUTLEDGE

Mixing Audio

Mixing Audio: Concepts, Practices, and Tools is a vital read for anyone wanting to succeed in the field of mixing—covering the entire mixing process, from fundamental concepts to advanced techniques.

Packed full of photos, graphs, diagrams, and audio samples, it teaches the importance of a mixing vision, how to craft and evaluate your mix, and then take it a step further. *Mixing Audio* describes the theory, the tools, and how these are put into practice while creating mixes. The companion website, featuring over 2,000 audio samples as well as Multitrack Audio Sessions, is a perfect complement to the book.

The new edition includes:

• An extension of the discussion on mastering which now includes the loudness war, LUFS targets, and DIY mastering
• Updated figures and illustrations throughout
• A new section covering console emulation

Roey Izhaki holds a BA in recording arts and has been mixing since 1992. As an audio engineering academic lecturer for 10 years, he has given mixing and audio seminars across Europe.

Mixing Audio

Concepts, Practices, and Tools

Fourth Edition

Roey Izhaki

Routledge
Taylor & Francis Group

Contents

Symbols used

Audio samples
www.MixingAudio.com/audio
Tracks referenced within these boxes are included on the website, organized in folders by chapter. Please mind your monitoring level when playing these tracks.

Notes
These boxes contain tips or other ideas worth remembering.

Introduction

It's not often a new form of art is conceived; where or when the art of mixing was born is not easy to answer. We can look at the instrumentation of orchestral pieces as a very primitive form of mixing—different instruments that played simultaneously could mask one another; composers understood this and took it into account. In the early days of recording, before multitrack recorders came about, a producer would place musicians in a room so that the final recording would make sense in terms of levels and depth. Equalizers, compressors, and reverbs hadn't yet been invented; there was no such role as a mixing engineer either, but sonically combining various instruments in order to produce an appealing, coherent, and aesthetic sound was an ambition shared by many.

Like many other new forms of creative expression that emerged in the twentieth century, mixing was tied to technology. It was the appearance of the multitrack tape machine during the 1960s that signified the dawn of mixing as we know it today. Yes, there was a time when having the ability to record eight instruments separately was a dream come true. Multitracks allowed us to repeatedly play recorded material before committing sonic treatment to the mix. Equalizers, compressors, and reverbs soon became customary in studios; audio consoles grew in size to accommodate more tracks and facilities. We had more sonic control over individual tracks and over the final master. The art of mixing was flourishing. Music sounded better.

The 1990s significantly reshaped much of the way music is made, produced, recorded, mixed, and even distributed—computers triumphed. Realtime audio plugins were first introduced with the release of Pro Tools III as far back as 1994, but such a setup required a dedicated DSP card. It was Steinberg's 1996 release of Cubase VST that gave us the audio plugins we now take for granted—a piece of software that can perform realtime audio calculations using the computer's CPU. The term *project studio* was soon coined as computers became more affordable and capable, and the hiring of expensive studios was no longer a requisite for multitracking and mixing. However, the processing power of computers back then could still not compete with the quality and quantity of mixing devices found in a professional studio.

Things have changed—running ten quality reverbs simultaneously on a modern DAW has been a reality for some time. There are now more audio plugins in the market than hardware units, and the quality of these plugins is constantly improving. Professional studios will always, it seems, have an advantage over project studios, if only for their acoustic qualities. However, DAWs offer outstanding value for money, constantly improving quality and widening possibilities.

1

So is everything rosy in the realm of mixing? Not quite. It is thanks to computers that mixing has moved from large and expensive studios into bedrooms. More people than ever are mixing music, but only a few can be labeled experts. Mixing used to be done by skilled engineers, who were familiar with their studio and the relatively small set of expensive devices it contained. Mixing was their occupation—and for many their *raison d'être*. On the contrary, project studio owners generally do much more than just mixing—for many, it is just another stage in an independent production chain. So how can these people improve their mixing, specifically when time is often constrained?

This is where this book comes in. When the first word of this book was typed back in 2004, mixing literature was limited, cluttered, and often only scratched the surface. This book was originally conceived to stand as a much-needed comprehensive source. It is hard to believe that nowadays, we have the opposite problem—there is too much out there, and, if literature isn't enough, there are also blog posts, online forums, and video tutorials. As we humans are increasingly and involuntarily assuming the role of information filters in this vast jungle called the web, a book such as this can spare many the foraging—what you need or wish to know in one place.

Being comprehensive doesn't come without a cost, though. This book is long; painfully long if you ask me. Perhaps of little comfort is that this book isn't quite a cover-to-cover type of read—feel free to stop reading now, look at the table of contents, and jump to the topic of most interest to you. Possibly not everything will be clear without reading some preceding chapters, but you should grasp the bulk of it. Regardless, many readers have testified that with so much to digest, it was only on the second or third reading understanding sunk in.

I would like, in this opening text, to expose the greatest misconception that exists about mixing: it is wrongly assumed by some that mixing is a purely technical service, and some even declare that mixing is simply a remedy for imperfect recordings. There is no doubt that mixing entails technical aspects: a problematic level balance, uncontrolled dynamics, and deficient frequency response are just a few of the technical issues we encounter. Yet, with the right amount of effort, almost anybody can master the technical aspects of mixing—after compressing 100 vocal tracks, one should be getting the hang of it. Technical skills are advantageous but can be equally acquired by all. The true essence of mixing does not lie in these skills. Many mixes are technically great, but nothing more than that; equally, many mixes exhibit some technical flaws, but as a listening experience they are breathtaking. It is for their sheer creativity—not for their technical brilliance—that some mixes are highly acclaimed and their creators deemed sonic visionaries.

The sonic qualities of music are inseparable from the music itself—the Motown sound, the Neve sound, the Wallace sound, and so forth. The nontechnical side of mixing entails crafting the sonic aspects of music: shaping sounds, crystallizing soundscapes, establishing harmony between instruments, and building impact—all rely on the many creative decisions that we make; all are down to the talent and vision of each individual; all have a profound influence on how the music is perceived. It is in the equalization we dial, in the reverb we choose, in the attack we set on the compressor, to name but a few. There simply isn't one correct way of doing things—be it an acoustic guitar, a kick, or any other instrument, it can be mixed in 100 ways; all could be considered technically correct, but some would be more remarkable than others. A mix is a sonic portrait of

the music. The same way different portraits of a person can each project a unique impression, different mixes can convey the essence of the music in extremely different ways. We are mixing engineers, but more importantly: we are *sonic artists*.

By the time you finish reading this book, you should have far more knowledge, a greater understanding, and improved auditory skills that will together enable you to craft better mixes. However, I hope that you keep this in mind:

Mixing is an art.

A friendly warning—mind your listening levels

It would not make sense for wine tasters to sip boiling oil, just as it would not make sense for mixing engineers to stick sharp needles into their eardrums. While I have yet to meet an engineer who fancies needles in his or her eardrums, very loud levels can be equally harmful. Unlike needle-sticking, the hearing damage caused by loud levels is often not immediate, whether involving short or long periods of exposure.

Sparing the medical terminology, with years one might lose the ability to hear high frequencies, and the really unlucky could lose substantial hearing ability. In some circumstances, very loud levels can cause permanent damage to the eardrum and even deafness. Most audio engineers have had one or two level-accidents; the majority of us are fine. But hearing a continuous 7 kHz tone is no laughing matter, especially when it lasts for three days.

The allowance, as they say in Italian, is *forte ma non troppo*—loud but not too much. The National Institute for Occupational Safety and Health in the USA recommends that sound exposure to 85 dBSPL should not exceed eight hours per day, halving the time for each 3 dB increase. A quick calculation reveals that it is only safe to listen to 100 dBSPL for 15 minutes. A screaming child a meter away is roughly 85 dBSPL. A subway train one meter away produces roughly 100 dBSPL when cruising at normal speed.

I have done my best to keep any audio samples at relatively consistent levels. Still, some samples had to be louder than others. Please mind your monitoring level when listening to these samples. Remember that too quiet can easily be made louder, but it might be too late to turn down levels once they are too loud.

Why we like loud levels so much is explained in Chapter 2. But if we are all to keep enjoying music, all we have to do is be sensible about the levels at which we mix and listen to music.

Levels, like alcohol, are best enjoyed responsibly.

Part I

Concepts and practices

1 Music and mixing

Music: An extremely short introduction

You love music. All of us are mixing because music is one of our greatest passions, if not *the* greatest. Whether starting as a songwriter, bedroom producer, performer, or studio tea boy, we were all introduced to mixing through our love of music and the desire to take part in its creation.

Modern technology dictates, to some extent at least, how we go about our life: we watch more and read less, we message more and talk less, we look at our smartphones more and less at one another. As far as music is concerned, new technologies have provided new opportunities, increased reach, and improved quality. The invention of the wax cylinder, radio transmission, tapes, CDs, software plugins, iTunes, smartphones, and Spotify has made music more readily accessible, widely consumed, and easier to create. One of mankind's most influential inventions—the Internet—is perhaps music today's greatest catalyst. Nowadays, a computer or a smartphone is all one needs to browse and listen to music. Music is universal and all-encompassing. It is in our living rooms, in our cars, in malls, on our televisions, and in hairdressing salons.

There is a strong bond between music and mixing, and to understand it we should start by discussing the not-too-distant past. Sacred music dominated the Western world until the nineteenth century, with most compositions commissioned for religious purposes. Secular music has evolved throughout the years, but really took off with the arrival of Beethoven—a daring and innovative composer at the time, no doubt. But it was the way that his music made people *feel* that changed the course of music so dramatically. Ernest Newman once wrote about Beethoven's symphonies:

> *The music unfolds itself with perfect freedom; but it is so heart-searching because we know all the time it runs along the quickest nerves of our life, our struggles & aspirations & sufferings & exaltations*.[1]

We can easily identify with this when we think about modern music—there is no doubt it can have a huge impact on us. Following Beethoven, music became a love affair between two willing individuals—the artist and the listener—fueled by what is today an inseparable part of music: emotions.

7

Today, music rarely fails to trigger emotions—all but a few pieces of music have some sort of mental or physical effect on us. "Killing in the Name" by Rage Against the Machine can trigger a feeling of rage or rebellious anger. Others find it hard to remain stationary when they hear "Hey Ya!" by OutKast. Music can turn a bad morning into a good one. Music can also trigger sad or happy memories, and so the same good morning can turn into a more retrospective afternoon after hearing Albinoni's "Adagio" (which goes to show that it's not just emotive lyrics that affect us).

As we shall soon see, our response to music mostly stems from our unconscious mind. Yet, we sometimes deliberately listen to music in order to incite a certain mood—some listen to ABBA as a warm-up for a night out, others to Iggy Pop. Motion-picture directors understand well how profoundly music can affect us and how it can be used to solicit certain emotional responses from the audience. We all know what kind of music to expect when a couple falls in love or when a shark is about to attack; it would be a particular genre of comedy that used "YMCA" during a funeral scene.

As mixing engineers, one of our prime functions, which is actually *our responsibility*, is to help deliver the emotional context of a musical piece. From the general mix plan to the smallest reverb nuances, the tools we use—and the way we use them—can all sharpen and even create power, aggression, softness, melancholy, psychedelia, and many other moods. Mostly, it would make little sense to distort the drums on a mellow love song, just as it would not be right to soften the beat of a hip-hop production. When approaching a new mix, we may ask ourselves:

• What is this song about?
• What emotions are involved?
• What message is the artist trying to convey?
• How can I support and enhance the song's vibe?
• How should the listener respond to this piece of music?

As basic as this idea might seem, it is imperative to comprehend—it is emotions that gel the music and mix together, not technical excellence.

A mix can, and should, enhance the music: its mood, the emotions it conveys,
and the response it should incite.

The role and importance of the mix

A basic definition of mixing is: A process in which multitrack material—whether recorded, sampled, or synthesized—is balanced, treated, and combined into a multichannel format (most commonly, two-channel stereo). But a less technical definition would be: **a sonic presentation of emotions, creative ideas, performance, and musicianship**.

Even for the layperson, sonic quality does matter. Take talking on a cell phone, for example—people find it annoying when background noise masks the other party. Intelligibility is the most elementary requirement when it comes to sonic quality, but it goes far beyond that. Today's cell phones are hardly better than playback systems from the 1950s. It is no wonder that people prefer listening to music via their kitchen's mini-

system or the living room hi-fi. What would be the point of more expensive hi-fi systems if the sound quality were no better than a cell phone speaker?

Sonic quality is also a powerful selling point. It was a major contributor to the rise of the CD and the fall of compact cassettes. Novice classical music listeners often favor new recordings over older, monophonic ones, regardless of how acclaimed the performance on these early recordings is. Record companies issue digitally remastered versions of classic albums that allegedly sound better than the originals. The once-ubiquitous iPod owed much of its popularity to the MP3 format—no other lossy compression format has managed to produce audio files so small, yet of an acceptable sonic quality.

> *The majority of people appreciate sonic quality more than they realize.*

It is our responsibility as mixing engineers to craft the sonic aspects of the final mix. This involves how different instruments combine, but also how each sounds—the whole and its parts. Let us consider for a moment the differences between studio and live recordings. During a live concert, there are no second chances. You are unable to rectify sloppy performance or a buzz from a faulty DI box. Both the recording equipment and the environment are inferior compared with those found in most studios—it would be unreason able to place Adele in front of a *U87* and a pop shield during a live show. Also, when a live recording is mixed on location, a smaller and cheaper arsenal of mixing equipment is used. All of these constraints could result in different instruments suffering from masking, poor definition, erratic dynamics, and deficient frequency response, to name just a few possible problems. Audio terms aside, these can translate into a barely audible bass guitar, honky lead vocals that come and go, a kick that lacks power, and cymbals that lack spark. Altogether, these can make a live recording less appealing. A studio recording is not immune to such problems, but in most cases it provides much better raw material to work with, and, in turn, better mixes. With all this in mind, the true art of mixing is far more than just making things sound right …

Many people are familiar with Kurt Cobain, Dave Grohl, and Krist Novoselic as the band members of Nirvana, who, back in 1991, changed the face of alternative rock with the release of *Nevermind*. The name Butch Vig might ring a bell for some, but the general public will be unlikely to have heard of Andy Wallace. The front cover of my *Kill Bill* DVD makes it extremely difficult to ignore Tarantino's writer and director credits. But it is seldom that an album cover credits the producer, let alone the mixing engineer. Arguably, the production of Dr. Dre can be just as important as the artists he produces, and perhaps *Nevermind* would have never been such an enormous success had it not been for Andy Wallace's mixes. Nevertheless, record labels generally see very little marketing potential in production personnel. Ironically, major record companies do part with large sums of cash in order to have a specific engineer mix an album because they all realize that:

> *The mix plays an enormous role in an album or track's success.*

To understand why, one may wish to compare Butch Vig's and Andy Wallace's mixes for Nirvana's "Smells Like Teen Spirit" (both can easily be found online through streaming services).

Both Vig and Wallace used the same raw tracks, yet their mixes are distinctly different. Vig's mix entails some unbalanced frequency spectrum that involves masking and the absence of spark; a few mixing elements, such as the snare reverb, are easily discernible. Wallace's mix is burnished and balanced; it boasts high definition and perfect separation between instruments; the ambiance is present, but like many mixing elements it is subtle.

Perhaps the most important difference between the two mixes is that Vig's mix sounds more natural (more like a live performance), while Wallace's mix sounds more artificial. It is not equipment, time spent, or magic tricks that made these two mixes so dissimilar—it is simply two different sonic visions. Vig has opted for the real and organic, whereas Wallace, a sonic alchemist who was perfecting his polishing skills at the time, combined every aspect of this powerful song into an extremely appealing masterpiece, albeit not a live-sounding one. Like many other listeners, Gary Gersh—Geffen Records' A&R—liked it better.

Straight after recording *Nevermind*, it was Vig that started mixing the album. Having spent countless hours in the recording studio listening to the same songs, it is common for a producer to wear out and develop sonic biases. A tight schedule and some artistic disagreements with Cobain left everyone (including Vig) feeling that it would be wise to bring fresh ears in to mix the album. From the bottom of the prospective engineers list, Cobain chose Wallace, mostly due to his mixing credits for Slayer. Despite Nirvana approving the mixes, following *Nevermind*'s extraordinary success, Cobain complained that the overall sound of *Nevermind* was too slick—perhaps suggesting that Wallace's mixes were too listener-friendly for his somewhat anarchic and unrefined taste. Artistic disagreements are something engineers come across often, especially if they ignore the artist's musical values. Yet, some suggested that Cobain's retroactive complaint was only a mis-targeted reaction to the massive success and sudden fame the album brought. Not only did *Nevermind* leave its mark on music history; it also left a mark on mixing history—its sonic legacy, a part of what is regarded as the *Wallace sound*, is still an aspiration for many today. As testament to Wallace's skill, *Nevermind* has aged incredibly well and still impresses despite enormous advances in mixing technology.

The example above is just one of many that demonstrate how a good mix can sharpen the emotional message of a musical piece, make it more appealing to the listener, and boost commercial success. Conversely, a bad mix can negatively affect a potentially great production and significantly impair its chance of success. This is not only relevant for commercial releases. The price and quality of today's DAWs enable unsigned artists and bedroom producers—given sufficient dedication—to craft mixes that are of an equal standard to commercial mixes. For quite some time now, A&Rs are receiving demos of a respectable mix quality. Just as a studio manager might filter through a pile of CVs and eliminate candidates based on poor presentation, an A&R might dismiss a demo for its poor mix.

Mixing engineers know what a dramatic effect mixing can have on the final product. With the right amount of effort, even the poorest recording can be made appealing. Yet, there are a few things we cannot do; for example: correct a truly bad performance, compensate for a very poor production, or alter musical ideas. If the piece does not

have potential to begin with, it will fail to impress the listener, no matter how note-worthy the mix is.

A mix is as good as the song.

The perfect mix

It doesn't take much experience before the novice mixer can begin to recognize problems in a mix. For instance, we quickly learn to identify vocals that are too quiet or a deficient frequency response. We will soon see that, once a mix is problem-free, there are still many things we can do in order to make it better. The key question is: What is better?

At this point, I recommend an exercise called **excerpt set** (Figure 1.1)—an essential mixing experiment. It takes around half an hour to prepare, but provides a vital mixing lesson. The excerpt set is very similar to a DJ set, except each track plays for around 20 seconds and you do not have to beat-match. Simply import around 20 songs from your music library into your audio sequencer. Then trim a random excerpt of 20 seconds from each track and arrange the excerpts consecutively. It is important to balance the perceived level of all excerpts, and cross-fade them. Now listen to your set, beginning to end, and notice the differences between the mixes; these differences are highly likely to be great. You might also learn that mixes you thought were good are not as good when played before or after another mix. While listening, try to note mixes that *you* think overpower others. This exercise will help develop a heightened awareness of what a good mix is and why.

Most of us do not have a permanent sonic standard stored in our brains, so a mix is only better or worse than the previously played mix. The very same mix can sound dull compared with one mix but bright compared with another. (With experience, we develop the ability to critically assess mixes without the need for a reference, although usually only in a familiar listening environment.) In addition, our auditory system has a very quick settle-in time, and it becomes accustomed to different sonic qualities so long as these remain constant for a while. In essence, all our senses work that way—a black-and-white scene in a color movie is more noticeable than the lack of color on a black-and-white TV. The reason why the excerpt set is such an excellent tool for revealing differences is that it does not give the brain a chance to settle in to a particular style. When mixes are played in quick succession, we can more easily perceive the sonic differences between them.

Different engineers have different ideas and mix in different environments, and therefore produce different mixes. Our ears are able to tolerate radical differences as long as mixes are not heard in quick succession. It is hard to find two albums that share an identical sound because different genres are mixed differently—jazz, grime, and trance will rarely share the same mixing philosophy; different songs involve different emotions and therefore call for different soundscapes; and the quality and nature of the raw tracks vary between projects. But we shouldn't forget that each mixing engineer is an artist in their own right, and each has different visions and ideas about what's best. Asking what is a perfect mix is like asking who is the best writer that ever lived, or who is the greatest basketball player of all time—it is subjective, situational, and highly contextual.

Figure 1.1 Excerpt set. This sequence of 20-second excerpts from various productions is used as an important comparison tool between mixes.

Mixing engineers will often adjust their style depending on the project. One example is Rich Costey, who mixed Muse's Absolution, *imbuing it with a very polished feel. He later produced Franz Ferdinand's* You Could Have It So Much Better, *taking a much rawer mixing approach with a distinctly retro feel. His mixes on Glasvegas's debut album are anthemic and sharp-sounding, featuring dominant reverbs typical of mixes from the 1980s.* Humbug *by Arctic Monkeys feels dark and beefy and features a contemporary sound with retro touches. Each of these different mixing approaches works a charm for its respective album.*

Note

1 Allis, Michael (2004). Elgar, Lytton, and The Piano Quintet, Op. 84. *Music & Letters*, Vol. 85 No. 2, pp. 198–238. Oxford University Press. Originally a letter from Newman to Elgar, January 30, 1919.

2 Some axioms and other gems

Louder is better

In 1933, two researchers at Bell Labs—Harvey Fletcher and W.A. Munson—conducted one of the most significant experiments in psychoacoustics. Their experiment was based on a series of tests taken by a group of listeners. Each test involved playing a test frequency followed by a reference tone of 1 kHz. The listener simply had to choose which of the two was louder. Successive tests involved either a different test frequency or different levels. Essentially, what Fletcher and Munson tried to conclude is how louder or softer different frequencies had to be in order to be perceived as loud as 1 kHz. They compiled their results and devised a chart known as the *Fletcher–Munson curves*. A chart based on the original Fletcher–Munson study is shown in Figure 2.1. I am presenting it upside down, as it bears a resemblance to the familiar frequency-response graphs that we see on some equalizers, with peaks at the top. A similar experiment was conducted two decades later by Robinson and Dadson (resulting in the *Robinson–Dadson contours*), and today we use the ISO 226 standard (which is still subject to occasional revisions). The formal name for the outcome of these studies is *equal-loudness contours*.

Each curve in Figure 2.1 is known as a *phon curve*, labeled by the level of the 1 kHz reference. To give an example of how to read this chart, we can follow the 20-phon curve and see that, if 1 kHz is played at 20 dBSPL, 100 Hz would need to be played at 50 dBSPL in order to appear equally loud (a 30 dB difference, which is by no means marginal). The graph also teaches us that our frequency perception has a bump around 3.5 kHz—this is due to the resonant frequency of our ear canal. Interestingly, this is pretty much the center frequency of a baby's cry.

One important thing that the equal-loudness contours teach us is that we are more sensitive to mid-frequencies—an outcome of the lows and highs roll-off that can be seen on the various curves. But more importantly, it is evident that at louder levels our frequency perception becomes more even—the 0-phon curve in Figure 2.1 is the least flat of all the curves; the 100-phon curve is the most even. Another way to look at this is that the louder music is played, the louder the lows and highs are perceived. In extremely general terms, we associate lows with power and highs with definition, clarity, and spark. So it is only natural that loud levels make music more appealing—louder is perceived as better.
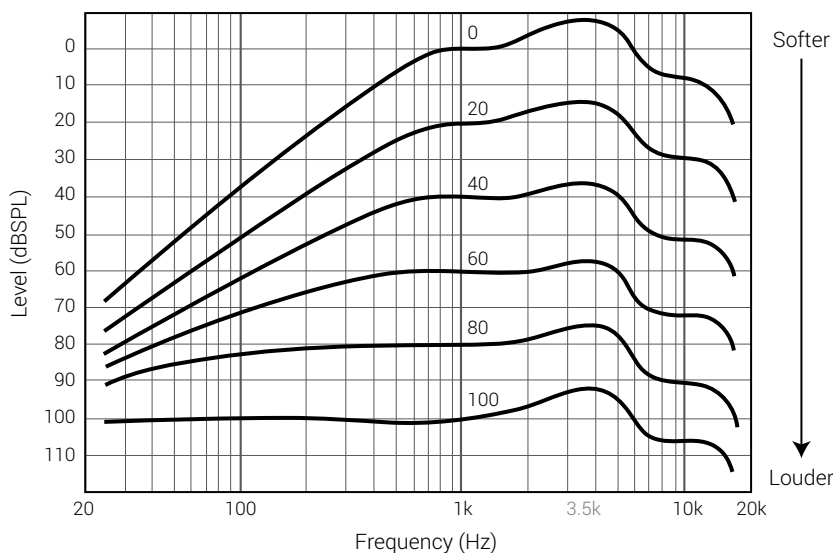
Figure 2.1 The Fletcher–Munson curves (shown here upside down). Note that on the level axis, soft levels are at the top, loud at the bottom.

This phenomenon explains the ever-rising level syndrome that many experience while mixing: once levels go up, it is no fun bringing them down. The more experienced among us develop the discipline to defeat this syndrome by keeping levels constant.

The louder music is played, the more lows and highs we perceive compared with mids.

The latest ISO 226 contours are slightly different than those shown in Figure 2.1*; they show an additional bump around 12 kHz and a steeper low-frequency roll-off, which also occurs on the louder phon curves.*

The fact that our frequency perception alters in relation to levels is a fundamental mixing issue. How are we supposed to craft a balanced mix if the frequency content varies with level? At what level should we mix? And what will happen when the listener plays the track at different levels? The answer is: we check our mix at different levels, and try to make it as level-proof as possible. We know what to expect when we listen at softer levels—less highs and lows. It is possible to equalize the different instruments so that even when the highs and lows are softened, the overall balance between instruments hardly changes. For example, if the kick's presence is based solely on low frequencies, it will be heard less at quiet levels, if at all. If we ensure that the kick is also present on the high-mids, it will be heard much better at quiet levels. Many believe that the mids, which vary little with level, are the key to a balanced mix, and if the lows and highs are crafted as an extension to the mids, a mix will exhibit more stable balance at different levels. Also, many agree that if a mix sounds good when at low levels, it is likely to sound good when played loud; the opposite is not always true. Another pointer is that we can sometimes guess the rough level at which the mix is likely to be played

(e.g., dance music is likely to be played louder than ambient), and so we can use that level as the main reference while mixing (as reference means occasionally—mixing at nightclub levels throughout is next to certain to damage your ears).

Two common adages: The mids are the key to a balanced mix at varying levels.
A mix that sounds good at quiet levels is likely to sound good at loud levels.

There is another reason why louder is perceived as better. When listening at soft levels, we hear more of the direct sound coming from the speakers and less of the sound reflected from the walls (i.e., the room response). Sound energy is absorbed, mostly as it encounters a surface. The small amount of energy our speakers emit at quiet levels is absorbed by walls to a degree that only a fraction of it reflects back to our ears. At louder levels, more energy is reflected and we start hearing more of the room response. As a consequence, the louder music is played, the more we hear the reflections coming from around us, which provides us with the appealing sensation that the music surrounds us. There is an experiment you can do to demonstrate this effect, which might be more apparent with eyes shut—play a mix at quiet levels through speakers and try to define the spatial boundary of the sound image. Most people will imagine a line, or a very short rectangle between the two speakers. As the music is made louder, the sound image grows, and at some point the two-dimensional rectangle turns into a vague surrounding sense.

When making individual instruments louder in the mix, we perceive them better. The core reason for this is masking—the ability of one sound to cover up another. More specifically, the frequency ranges of one instrument mask those of another. One of the principal rules of masking is that louder sounds overpower quieter sounds. The higher the level of an instrument in the mix, the more it will tend to mask other instruments, and the more clearly it will be perceived. When focusing on a particular instrument, it is tempting to bring it up to hear it better, and once up it is likely to sound better. This may be misleading and lead to suboptimal balance.

Percussives weigh less

It is important to distinguish the different natures of the instruments we are mixing. An important mix resource is space; when different instruments are combined, they compete for that space (mostly due to masking). Percussive instruments come and go—a kick, for example, has little to no sound between various hits. Percussives fight for space in successive, time-limited periods. On the other hand, sustain instruments play over longer periods and thus constantly fight for space. To give one extreme example, think of a rich pad produced using sawtooths (the most harmonically rich orderly waveform) that involves unison (an effect that spreads copies across the stereo image), and played in a legato fashion (long notes). Such a pad would fill both the frequency spectrum and the stereo panorama in a way that is most likely to mask many other elements in the mix.

In a practical sense, sustained instruments require somewhat more attention. Whether we are setting levels, panning, or equalizing them, our actions will have an effect over a longer period. Raising the level of a dense pad is likely to cause more masking problems than raising the level of a kick. If the kick masks the pad, it would only do so for

short periods—perhaps not such a big deal. But if the pad masks the kick, it would do so constantly—a big deal indeed.

Importance

Imagine yourself being on a *Seinfeld* set. In the scene being shot, Jerry and Kramer stand in a long line of people at a box office, engaged in conversation. Being the stars of the show, among all people, the production efforts would have been focused on the two stars. The makeup artist, for example, probably spent quite some time with them, perhaps little time with the extras standing next to them, and most likely no time with any other extras standing farther away in the line. In the camera shot, Jerry and Kramer are seen clearly in the center and extras are out of focus. The importance of the stars will also have been evident in the work of the gaffer, the grips, the boom operator, or any other crew member, even the chef.

Equally, different mix elements have varying importance within the mix. The importance of each instrument depends on many factors, like the nature of the production being mixed. In hip-hop, for example, the beat and vocals are generally the most important elements. In jazz, the snare is more important than the kick. Spatial effects are an important part of ambient music. A prominent kick is central to club music, but of far less importance in most folk music. Many more examples can be given. We also have to consider the nature of each instrument and its role in the overall musical context. Vocals, for example, are often of prime importance, but the actual lyrics also play a crucial role. The lyrics of Frank Sinatra's "My Way" are vital to the song's impact, and mixing a vocal part as such calls for more emphasis. Arguably, the lyrics to "Give It Away" by Red Hot Chili Peppers are of little importance to the overall song climate.

Importance affects how we mix different elements, be it levels, frequencies, panning, or depth we are working on. We will shortly look at how the order in which we mix different instruments and sections may also be affected. Identifying importance can make the mixing process all the more effective as it minimizes the likelihood of delving into unnecessary or less important tasks—for example, spending a fair amount of time on treating pads that only play for a short period of time at relatively low level. Those of us who mix under time constraints have to prioritize our tasks. In extreme circumstances, you might have as little as one hour to mix the drums, just half an hour for the vocals, and so on.

A useful question: How important is it?

Natural vs. artificial

A specific event that took place back in 1947 changed the course of music production forever. Patti Page, then an unknown singer, arrived at a studio to record a song called "Confess." The studio was set up in the standard way for that era, with all the performers in the same room, waiting to cut the song live. But there was a problem— "Confess" was a duet where two voices overlap, but for a reason yet to be found, no second vocalist showed up. Jack Rael, Page's manager, came up with the unthinkable:

Patti could sing the second voice as well, provided the engineer could find a way to *overdub* her voice. Legend has it that at that point, the engineer cried in horror: in real life, no person can sing two voices at the very same time. It's ridiculous. Unnatural! But to the A&R guy from Mercury Records, this seemed like a great gimmick that could secure a hit. To achieve this, the engineer did something that was never done before—cloning the track from one machine to another while adding the second voice on top. What then seemed so bizarre is today an integral part of music production.

For our purposes, a "natural" sound is one that emanates from an instrument that is played in our presence. If there are any deficiencies with the raw recordings (which capture the natural sound), various mixing tools can be employed to make instruments sound "more natural." A mix is considered more natural if it presents a realistic sound stage (among other natural characteristics). If natural is our goal, it would make no sense to position the kick up front and the rest of the drum kit behind it.

However, natural is not always best—natural can also be seen as very ordinary. Early on in photography, it occurred to people that shadows, despite being such a natural part of our daily life, impair visuals. Most advertisements have had tone and color enhancements in order to make them look "better than life." The same goes for studio recording. It is not uncommon today to place the kick in front of the drum kit, despite the fact that this creates a very unnatural spatial arrangement.

One of the principal decisions we make when we began a mix is whether we want things to sound natural or artificial. This applies on both the mix and instrument levels. Some mixes call for a more natural approach. Jazz enthusiasts, for example, expect a natural sound stage and natural-sounding instruments, although in recent years more and more jazz mixes involve an unnatural approach—for instance, compressed drums with an emphasized kick and snare. This fresh, contemporary sound has attracted a new audience (and even some connoisseurs), and facilitated a wider market for record companies to exploit. Popular music nowadays tends to be all but natural—the use of heavy compression, distortions, aggressive filtering, artificial reverbs, delays, distorted spatial images, and the like is routine. These paradigms, while not natural, increase the potential for creativity and profoundly affect the overall sound. Mixes are sonic illusions. The same way that color enhancement improves visuals, our mixing tools allow us to craft illusions that sound better or just different from real life. People who buy live albums expect a natural sound. Those who buy studio albums expect, to some extent, a sonic illusion, even if they don't always realize that.

Some inexperienced engineers are hesitant to process since they consider the raw recording a natural touchstone. Often they are cautious about even gentle processing, considering it to be harmful. Listening to a commercial track that was mixed with an artificial approach will reveal just how extreme mixing treatments can be. Take vocals, for example: their body might be removed, they might be compressed so that there are no dynamic variations, or they might be overtly distorted. We have to remember that radical mixing is generally unperceived by those without a trained ear—the majority of listeners, that is. Here are three sentences my mother has never said and will probably never say:

- Listen to her voice. It's over-compressed.
- That guitar is missing body.
- The snare is too loud.

The common listener does not think or speak using these terms. For them, it is either exciting or boring; they either feel it or don't; and most importantly, they either like it or they don't. This leaves a lot of room for wild and adventurous mixing treatments—we can filter the hell out of a guitar's bottom end; people will not notice. We can make a snare sound like a Bruce Lee punch; people will not notice. Just to prove a point here, the verse kick on Nirvana's "Smells Like Teen Spirit" reminds me more of a bouncing basketball than any bass drum I have ever heard playing in front of me. People do not notice.

3 Learning to mix

An analogy can be made between the process of learning a new language and that of learning to mix. At the beginning, nothing seems to make sense. With language, you are unable to understand simple sentences or even separate the words within a sentence. Similarly, if you play a mix to most people, they will not be able to hear a reverb or compression as they haven't focused on these sonic aspects before, let alone used reverbs or compressors. After learning some individual words and how to use them, you find yourself able to identify them in a sentence; in the same way, you start learning how to use compressors and reverbs and then learn to recognize them in mixes. Pronouncing a new word can be challenging, since it is not easy to notice the subtle pronunciation differences in a new language, but after hearing and repeating a word 20 times, you get it right; likewise, after compressing 20 vocal tracks, you will start to identify degrees of compression and quickly evaluate what compression is most suitable. Then you will begin to learn grammar so that you can begin to connect words together and construct coherent sentences, much like all your mixing techniques help you to craft a mix as a whole. Finally, since conversation involves more than one sentence, the richer your vocabulary is and the stronger your grammar, the more sentences you are able to properly construct. In mixing, the more techniques and tools you learn and the more mixes you craft, the better your mixing becomes. Practice makes perfect.

What makes a great mixing engineer?

World-class mixing engineers might earn more for a single album than many people earn in a year. Some mixing engineers also receive points—a percentage from album sale revenue. On both sides of the Atlantic, an accomplished mixing engineer can enjoy a six-digit annual revenue. These individuals are being remunerated for their knowledge, experience, and skill. Record labels reward them for that, and in exchange enjoy greater sales.

It is clear why mixing is often done by a specialized person. And it is such a vast area that it is no wonder some people devote themselves entirely to it—the amount of *knowledge* and *practice* required to make a great mixing engineer is vast.

Primarily, the creative part of mixing revolves around the three steps shown in Figure 3.1. The ability to progress through these steps can lead to an outstanding mix.

How do I want it to sound?

Vision

Evaluation   →   Action

Does it sound like I want it to?   What equipment should I use?
Does it sound right?   How should I use the equipment?
What is wrong with it?
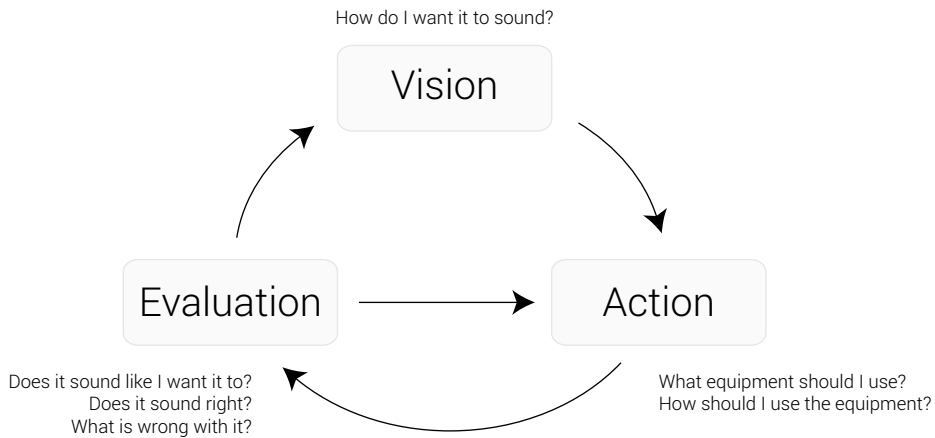
Figure 3.1 The three steps of creative mixing.

But a great mixing engineer will need a notch more than that, especially if hired. These steps are explained, along with the requisite qualities that make a great mixing engineer, in the following sections.

Mixing vision

There are different methods of composing. One involves using an instrument, say a piano, and then, either by means of trial and error or using music theory, coming up with a chord structure and melody lines. Another approach involves imagining or thinking of a specific chord or melody and then playing it. The latter process of "visualizing" and then playing or writing is favored by many composers and songwriters.

The same two approaches apply to mixing as well. If we take the equalization process of a snare, for example, the first approach involves sweeping through the frequencies, then choosing whatever frequency appeals to us most. The second approach involves first imagining the desired sound and then approaching the EQ in order to attain it. Put another way, the first approach might involve thinking, "OK, let's try to boost this frequency and see what happens," while the second might sound more like, "I can imagine the snare having less body and sounding more crisp." Just as some composers can imagine the music before they hear it, a mixing engineer can imagine sounds before attaining them—a big part of mixing vision. Mixing vision is primarily concerned with the fundamental question: **How do I want it to sound?** The answer could be soft, powerful, clean, etc. But mixing vision cannot be defined by words alone—it is a sonic visualization, which later manifests through the process of mixing.

The selection of tools that we can use to alter and embellish sound is massive—equalizing, compressing, gating, distorting, adding reverb or chorus are just a few. So what type of treatment should we use? There are infinite options available to us within each category—the frequency, gain, and Q controls on a parametric equalizer provide billions of possible combinations. So why should we choose a specific combination and not another? Surely equalizing something in a way that makes it sound right does not mean that a different equal ization would not make it sound better. A mixing vision provides the answer to these questions: "because this is how I imagined it; this is how I wanted it to sound."

Evaluation ⟶ Action

Does it sound right?   What equipment should I use?
What is wrong with it?   How should I use the equipment?
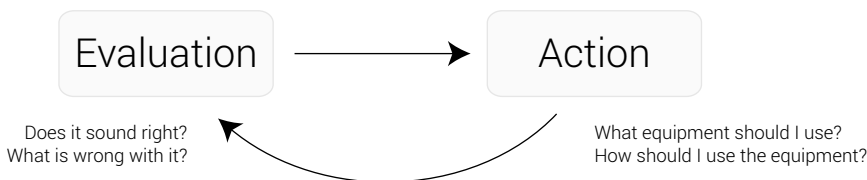
Figure 3.2 The novice approach, without a mixing vision.

A novice engineer might lack imagination. The process of mixing for him or her is a trial-and-error affair between acting and evaluating (Figure 3.2). But how can one critically evaluate something without a clear idea of what one wants in the first place? Having no mixing vision can make mixing a very frustrating hit-or-miss process.

Having a mixing vision can make all the difference between the novice and the professional mixing engineer. While the novice shapes the sounds by trial and error, the professional imagines sounds and then achieves them.

The skill to evaluate sounds

The ability to craft a good mix is based on repeated **evaluations**. One basic question, often asked at the beginning of the mixing process, is **"What's wrong with it?"** Possible answers might be, "the highs on the cymbals are harsh," "the frequency spectrum of the mix is too heavy on the mids," or "the drums are not powerful enough." From the endless amount of treatment possibilities we have, focusing on rectifying the wrongs provides a good starting point. It can also prevent the novice from doing things that aren't actually necessary (for example, equalizing something that didn't really require equalization) and thereby save precious studio time.

At times, it might be hard to tell what is wrong with the mix, in which case our mixing vision provides the basis for our actions. After applying a specific treatment, the novice might ask, **"Does it sound right?"** while the veteran might also ask, **"Does it sound the way I want it to?"** Clearly, the veteran has an advantage here since this question is less abstract.

Mastering the tools and knowledge of other common tools

Whether with or without a clear mixing vision, we perform many **actions** in order to alter sounds. When choosing a reverb for vocals, the novice might tirelessly go through all the available presets on a reverb emulator. There can be upwards of 50 of these, and so the entire process can take some time. The veteran, on the other hand, will probably quickly access a specific emulator and choose a familiar preset; a bit of tweaking and the task is done. It takes very little time. Experienced mixing engineers know, or can very quickly find out, which tool will do the best job in a specific situation. They can quickly answer the question: **"What equipment should I use?"**

Professional mixing engineers do not always work in their native environment. They will sometimes work in different studios and, even though they might take their favorite gear with them, a big part of the mix will be done using in-house equipment.

Therefore, professional mixing engineers have to be familiar with the common tools found in a commercial environment.

Mastering the tools at one's disposal does not only mean having the ability to pick the right tool for a specific task, but also having the expertise to employ the equipment in the best way (**"How should I use the equipment?"**). Knowing whether to choose high-shelving or high-pass characteristics on an equalizer, or knowing that a specific compressor will work well on drums when more than one ratio button is pressed, are just a couple of examples.

It is also worth discussing the quantity of tools we have at our disposal. Nowadays, DAW users have a wider selection than those mixing using hardware. Not only are plugins cheaper, but they can be used across various tracks, simultaneously, whereas a specific hardware processor cannot. In an analog studio, a mixing engineer might have around three favorite compressors to choose from when processing vocals; DAW users might have a choice of 10. Learning each of these compressors—*understanding* each of them—takes time. Just reading the manual is time-consuming. Having many tools can mean that they are not realizing their potential because there is no time to learn and properly experiment with them all. Mixing is a simple process that only requires a pair of trained ears and a few quality tools. Less can be more.

Jack of all trades, master of none.

Theoretical knowledge

Four questions:

- When clipping shows on the master track in an audio sequencer, is it the master fader or all of the channel faders that should be brought down?
- To achieve more realistic results, should one or many reverb emulators be used?
- Why and when should stereo linking be engaged on a compressor?
- When should dither be applied?

To say that every mixing engineer knows the answers to these questions would be naïve. The same could be said that one cannot craft an outstanding mix without strong theoretical knowledge. There are more than a few highly successful engineers who would be unable to provide answers to many theoretical questions relating to their field. But knowing the answers to these questions is definitely an advantage. Knowledge is always a blessing and, in such a competitive field, can make all the difference. Out of two equally talented engineers with different levels of knowledge, it would not be difficult to choose who to work with.

To acquire knowledge, some might undertake an educational program while others might learn little by little on the job; but, either way, all mixing enthusiasts need to be compulsive learners—if the ratio on a compressor is set to 1:1, a novice will spend hours trying to figure out why no other control has an effect. Learning the difference between shelving and a pass filter is handy. The effect that dither has on the final mix quality is worth knowing. It would seem unreasonable for a mastering engineer not to know when to apply dither, but mixing engineers should know too.

It is better to know what you can do, and how to do it, than to understand what you have done.

Interpersonal skills

Studio producers need an enormous capacity to interact and deal with many people, with different abilities, moods, and degrees of dedication. Mixing engineers tend to work on their own and only occasionally mix in front of the client—whether that is the artist, A&R, or the producer. So, although to a lesser extent than a studio producer, like any job that involves interaction with people, mixing also requires good interpersonal skills.

When the band comes to listen to the mix, it should not come as a surprise that each band member will insist that his or her instrument is not loud enough. (In their defense, they are used to their instrument appearing louder to them, whether onstage or through the cans in the live room.) Even the old tricks of limiting the mix or blasting the full-range speakers will not always appease them. On many occasions, artists and A&R remark on the work of mixing engineers with the same rationale of accountants commenting on the work of graphic designers they have hired.

While the feedback from fresh ears can sometimes be surprisingly constructive, at times the clients' comments are either technically or artistically naïve or inappropriate. Things can easily become personal—mixing engineers, like the artists they mix, can become extremely protective about their work. Interpersonal skills can help avoid or resolve artistic disagreements, and assist with calmly expressing an opinion. But if the artist does not like some aspect of the mix, even if it's technically earth-shattering, however much the mixing engineer might disagree, it is the mixing engineer who must compromise. The client-is-always-right law is the same in mixing—after all, a displeased client is a lost client. For artists, each single or album goes on their CV forever and can be life-defining. Do you really want them to be unhappy with your work?

The ability to work fast

Learning something new can be tricky and testing—all guitar players experience some frustration before they can change chords quickly enough or produce a clean sound. It is maddening working on a single verse for a whole day and still being unhappy with the mix. But, as experience accumulates, it takes less time to choose tools and utilize them to achieve the desired sound. Also, our mixing visions become sharper and we can crystallize them more quickly. Altogether, each task takes less time, which leaves more time to elevate the mix or experiment. Needless to say, the ability to work fast is essential for hired mixing engineers, who work under busy schedules and strict deadlines.

Methods of learning

Reading about mixing

Literature is great. Books, magazine articles, and Internet forums can be the source of some extremely valuable theory, concepts, ideas, and tips. But reading about mixing

will not make a great mixing engineer, in the same way as reading a cookery book will not make a great chef. Reading about mixing gives us a better chance to understand core concepts and operate our tools, but the one thing it does not do is improve our sonic skills.

Reading manuals is also important, although unfortunately many people choose to neglect it. The basic aim of a manual is to teach us how to use our equipment, and sometimes also how to use it correctly or how to use it better. In their manuals, many manufacturers will present some practical advice on their products and sometimes on mixing in general. Sometimes the controls of a certain tool are not straightforward and it might take an eternity to understand what its function is without reading the manual.

Read the manual.

Reading and hearing

This book is an example of this method. An aural demonstration of mixing-related issues provides a chance to develop critical evaluation skills and better understanding of sonic concepts. While this method can contribute to all stages of mixing—vision, action, and evaluation—it is a passive way of learning since it does not involve active mixing.

Seeing and hearing

Watching other people mix is another way to learn. Many people want to work in a studio so they can learn from the experienced. Listening to others while they apply their expertise to a mix is a great opportunity and a valuable experience, but comes with two cautions. First, it is impossible to enter other people's minds—while watching them mix it might be possible to understand what they are doing, but not why they are doing it. **Mixing vision and experience are nontransferable**. Second, if we take into account the tricks and tips already published, what is left to learn from these experienced people is mostly their own unique techniques. True, learning the secret techniques of mixing engineers at the top of their game is great, but only if these are used in the right context. There is the belief that the greatest mixing engineers produce incredible mixes because of secret techniques. In practice, these amazing mixes are not down to secret techniques, but to an **extensive understanding of basic techniques and experience using them**. Secret techniques often only add a degree of polish or the individual's idiosyncratic sonic stamp.

Doing it

Without a shadow of a doubt, the best way to learn mixing is simply by doing it. Most of the critical skills and qualities of a great mixing engineer can be acquired through the practice of the art. While mixing, we learn to evaluate sounds and devices, use our equipment in the best way, work faster, and articulate our mixing vision quicker. Combined with good theoretical background and enough application, there is very little to stop anyone from becoming a competent mixing engineer. There is a direct link between mixing-miles and the final quality of the mix.

The best way to learn mixing is to mix.

Mixing analysis

Sometimes, learning the techniques of an art makes it hard to perceive the art as a whole. For example, while watching a movie, film students will analyze camera movements, lighting, edits, lip-sync, or acting skills. It can be hard for those students to stop analyzing and just enjoy movies like they did when they were fascinated kids. However, many mixing engineers find it easy to switch in and out from a mixing analysis state—even after many years of mixing, they still find it possible to listen to a musical piece without calculating how long the reverb is, where the trumpet is panned to, or questioning the sound of the kick. Others simply cannot help it.

Although it is far less enjoyable to analyze the technical aspects of a movie while watching it, this critical awareness can help make film students more conscientious filmmakers. Sit, watch, and learn how the masters did it—simple. The same approach works for mixing. Every single mix out there, whether good or bad, is a lesson in mixing. Learning is just a matter of pressing play and actively listening to what has been done. Although mixing analysis cannot always reveal how things were done, it can reveal much of what was done.

Your music collection contains hundreds of mixing lessons.

There are endless things to listen for when analyzing others' mixes, and these can cover any and every aspect of the mix. Here are just a few questions you might ask yourself while listening:

- How loud are the instruments in relation to one another?
- How are the instruments panned?
- How do the different instruments appear in the frequency spectrum?
- How far apart are the instruments in the depth field?
- How much compression was applied to the different instruments?
- Can any automation be detected?
- How long are the reverbs?
- How defined are the instruments?
- How do different mix aspects change as the song advances?

A quick demonstration seems appropriate here. The following points provide a partial mixing analysis for the first 30 seconds of Nirvana's "Smells Like Teen Spirit," the album version:

- The tail of the reverb on the crunchy guitar is audible straight after the first chord (0:01).
- There is extraneous guitar noise coming from the right channel just before the drums are introduced (0:05).
- The crunchy guitar dives in level when the drums are introduced (0:07).
- Along with the power guitars (0:09–0:25), the kick on the downbeats is louder than all other hits. (It appears to be the actual performance, but it can also be achieved artificially during mixdown.)
- When listening in mono, the power guitars lose some highs (0:09–0:25).
- The snare reverb changes twice (a particular reverb before 0:09, then no audible reverb until 0:25, then another reverb).

- During the verse, all the kicks have the same timbre (suggesting drum triggers).
- There is kick reverb during the verse.
- It is possible to hear a left/right delay on the hi-hats—especially during open/close hits. (This could be the outcome of a spaced microphone technique, but can also occur during mixdown.)
- The drums are panned audience-view.

The excerpt set (from Chapter 1) can be a true asset when it comes to mixing analysis, as the quick changes from one mix to another make many aspects more noticeable. Not every aspect of the mix is easily discernible: some are subliminal and are felt rather than heard. To be sure, the more time and practice we put into mixing analysis, the more we discover.

In addition to what we can hear from the plain mix, it is also possible to use different tools in order to reveal extra information. Muting one channel of the mix can disclose additional stereo information (e.g., a mono reverb panned to one extreme). Using a pass filter can help in understanding how things have been equalized. To reveal various stereo effects, one can listen in mono while phase-reversing one channel (this results in a mono version of the difference between the left and right, which tends to make reverbs and room ambiance very obvious).

Reference tracks

Mixing analysis is great, but it is impossible to learn hundreds of mixes thoroughly, and it can be impractical to carry them around just in case we need to refer to them. It is better to focus on a few select mixes, learn them inside out, analyze them scrupulously, and have them readily accessible.

Some mixing engineers carry a few reference tracks (mostly their own past mixes) so they can refer to them. The novice might refer to their reference tracks on a frequent basis. When mixing at home or in their studio, some have a specific folder on the hard drive with their select mixes.

In addition to reference tracks, including the excerpt set can be great since it enables a quick comparison between many different mixes. It is also possible to include a few raw tracks, which can later be used to evaluate different tools.

Our choice of reference tracks might not be suitable for every mix. If we are working on a mix that includes strings, and none of our reference tracks involve strings, it would be wise to look for a good mix that does. Likewise, if our reference tracks are all heavy metal and we happen to work on a chill-out production, it would be sensible to refer to some more appropriate mixes.

Usage of reference tracks

Reference tracks can be employed for different purposes:

- **As a source for imitation**—painting students often go to a museum to copy a familiar painting. While doing so, they learn the finest techniques of famous painters. Imitating another's techniques is part of the learning process. Likewise, there is nothing amiss in imitating proven mixing techniques—if you like the sound of the kick in a specific mix, why not imitate that sound in your mix? There is no reason why you can't replicate the technique of a specific track that you particularly like. When we are short of a mixing vision, we can replace it with the sonic image of an existing mix, or try to imitate it, or just some aspects of it. Trying to imitate the sound of a known mix is actually a great mixing exercise. However, caution must be exercised for several reasons. First, because productions can be so diverse that, whether in their emotional message, style, arrangement, quality, or nature of the raw material, what sounds good in another mix might not sound so good in yours. Second, setting a specific sound as an objective can be limiting and mean that nothing better will be achieved. Third, it is hard to obtain the same sounds with different recordings—when the ingredients are different, it's hard to make the dish taste similar. Finally, and most importantly, *imitation is innovation's greatest enemy*—there is little creativity involved in imitation. In fact, it might restrain the development of creative mixing skills.
- **As a source of inspiration**—while imitating a mix requires a constant comparison between the reference track and our own mix, reference tracks can be played before mixing to inspire us as to the direction in which the mix should go and what qualities it should incorporate. For the novice, such a practice can kick-start some mixing vision and set certain sonic objectives.
- **As an escape from a creative dead end**—sometimes we reach a point where we are clearly unhappy with our mix, but frustratingly cannot tell quite what is wrong with it. We might be simply out of ideas or lacking vision. Learning the difference between our mix and a specific reference mix can trigger new ideas, or suggest problems in our mix.
- **As a reference for a finished mix**—when we finish mixing, we can compare our mix to a specific reference track. Listening to how the professionals do it can help us generate ideas for improvement. The frequency response or relative levels of the two mixes are just two possible aspects that we might compare.
- **To calibrate our ears to different listening environments**—working anywhere but in our usual listening environment reduces our ability to evaluate what we hear. Just before we start to listen critically in an unfamiliar environment, whether mixing or just evaluating our own mixes, playing a mix we know well can help calibrate our ears to unfamiliar monitors, the acoustics, or even a different position within the same room.
- **To evaluate monitor models before purchase**—studio monitor retailers usually play customers a popular track that has an impressive mix, at loud levels. Chances are that the monitors will impress the listener that way. Listening to a mix that you are very familiar with can improve judgment.

It is worth remembering that if a reference track has been mastered, it is very likely to contain tighter dynamics, usually in the form of more allied relative levels and heavier compression. Also, mastered tracks are typically louder, whereas the overall loudness of a track is not a concern during the mixing stage. In some albums, frequency treatment takes place in order to match the overall sound to that of the worst track. These points are worth bearing in mind when comparing a reference track to a mix-in-progress—a mastered reference track is an altered version of a mix, usually for the better.

How to choose a reference track

Choosing some of our own past mixes for reference purposes is always a good idea. Having worked on these mixes, we are familiar with the finer details and, retrospectively, their faults. Ideally, reference materials should be a combination of both unmastered and commercial tracks. Here are a few of the qualities that reference tracks should have:

- **A good mix**—while subjective, your opinion of what is a good mix is central. It is important to choose a mix you like, not a production you like. Despite Elvis Presley's greatness, the sonic quality of his original albums is nowhere near today's standards.
- **A contemporary mix**—mixing has evolved. A good mix from the 1980s is likely to have more profound reverbs than the mix of a similar production from the 1990s. Part of the game is keeping up with the changing trends.
- **Genre related**—clearly, it makes little sense to choose a reference track of a genre that is fundamentally different from the genres you will be working on.
- **A dynamic production**—choosing a dynamic production, which has a dynamic arrangement and mix, can be like having three songs in one track. There is more to learn from such a production.

Reference tracks should not be:

- **A characteristic mix**—the mixing style of some bands, The Strokes, for example, is rather unique. A mix that has a distinct character will only serve those distinct productions and bands.
- **Too busy**—it is usually easier to discern mixing aspects in sparse productions.
- **Too simple**—the more there is to learn from a mix, the better. An arrangement made of a singer and their acoustic guitar might sound great, but will not teach you how to mix drums.

4 The process of mixing

The arrangement

The arrangement (or instrumentation) largely determines which instruments play, when, and how. An arrangement that lends itself to the mix will have its frequency, time, and stereo domains reasonably shared between its featured instruments.

Consider an arrangement that involves only one percussion instrument—a shaker. If panned center in a busy mix, it is most likely to be masked by other instruments; but panning it to one side can create an imbalanced stereo image. It might be easier for the mixing engineer to have a second percussion instrument, say a tambourine, so the two can be panned left and right. Then, if the shaker and tambourine are playing the same part, they will fight for the same space on the time and frequency domains, whereas if playing complementary parts, each should be more discernible.

One of the most relevant factors of the arrangement is its density. A sparse arrangement (Figure 4.1a) might call for a mix that fills various gaps in the frequency, stereo, and time domains. For example, an arrangement based solely on two tracks, an acoustic guitar and a vocal, could require creating something out of very little. At the other extreme, the challenge with a busy arrangement (Figure 4.1b) is often to create a

(a)   (b)
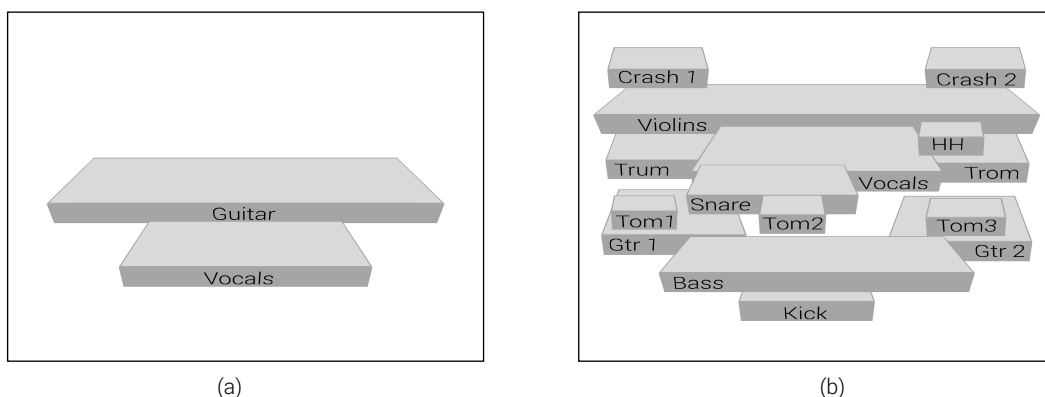
Figure 4.1 Sparse vs. dense arrangement.

space in the mix for each instrument. The busier the arrangement the harder it is to lay emphasis on a specific instrument or emphasize fine details.

Both Andy Wallace and Nigel Godrich faced sparse arrangements consisting of a guitar and vocal only in sections of "Polly" by Nirvana and "Exit Music" by Radiohead. Each tackled it in a different way—Wallace chose a plain, intimate mix, with fairly dry vocal and a subtle stereo enhancement for the guitar. Godrich chose to use very dominant reverbs on both the guitar and vocal. It is interesting to note that Wallace chose the latter, reverberant approach on his inspiring mix for "Hallelujah" by Jeff Buckley—an almost seven-minute song with an electric guitar and a single vocal track.

It is not uncommon for the final multitrack to include extra instrumentation along with takes that were recorded as try-outs, or in order to give some choices during mixdown. It is possible, for example, for the multitrack to include eight power-guitar overdubs; sometimes only mixing two of these tracks serves the mix best. Ditto for five different microphones that capture a kick drum.

There are opposite situations where the arrangement is so minimalist that it is very hard to produce a rich, dynamic mix. In such cases, nothing should stop the mixing engineer from adding instruments to the mix—so long time, talent, and ability allow this, and the client approves the additions.

It is OK to remove or add to the arrangement during mixdown.

It is worth remembering that the core process of mixing involves both **alteration** and **addition** of sounds—a reverb, for example, is an additional sound that occupies space in the frequency, stereo, and time domains. Some producers take this well into account by "leaving a place for the mix"—the famous vocal echo on Pink Floyd's "Us and Them" is a good example of this.

The recordings

Mixing in the production chain

Figure 4.2 illustrates the common production chain for recorded music. Producers may give input at each stage, but they often focus on the arrangement and recording stages. Each stage has an impact on the subsequent stage, yet it might be carried out by different people.
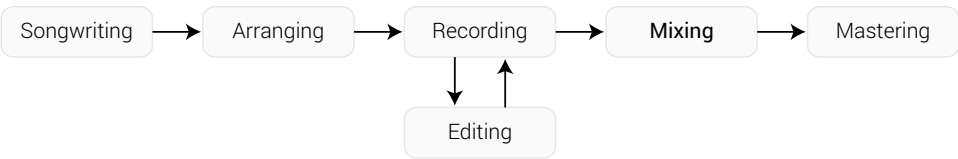
Figure 4.2 Common production chain for recorded music.

Mixing is largely dependent on both the arrangement and recording stages. A wrongly placed microphone can result in an acoustic guitar lacking body, which could be hard to rectify during mixdown. Some recording decisions are undoubtably mixing decisions. For example, the choice of stereo-miking technique for drum overheads determines the localization and depth of the various drums in the final mix. Altering these aspects during mixdown takes effort.

Mixing engineers, when a separate entity in the production chain, commonly face arrangement or recording issues such as those mentioned above. There is such a strong link between the arrangement, recordings, and mix that it is dubious for a producer or a recording engineer to have no mixing experience whatsoever. A good producer anticipates the mix. There is an advantage to having a single person helping with the arrangement, observing the recording process, and mixing the production. This ensures that the mix is borne in mind throughout the production process.

Recording the whole versus recording its parts

The production process has generally become more segregated through the years. In the early days of the recording industry, orchestras were placed in front of a single microphone and recorded into a monophonic medium. There was no mixing other than "physical mixing"—placing musicians in space to get the balance right. The practice of having all the musicians playing together continued even when multitracks appeared, and for high-budget projects this still happens today.

As technology evolved, more and more mixing was taking place—first to tape, then from tape—and the option to defer sonic decisions to a later stage became progressively attractive and economical. Singers would be placed in vocal booths, while instruments would be close-miked. To gain even further sonic control, musicians would stop playing together and instead take turns in the studio. In today's budget and home studios, this is pretty much the standard.

The paradigm used during the recording stage affects mixing a great deal. In the past and in many modern-day high-budget projects, the sonic aspects of the musical piece are shaped in the recording studio and great attention was given to the whole (how instruments sound together, which will in turn inform microphone placement and processing). This would normally leave the mixing engineer with the job of "spicing up" a production that is already fairly mixed.

In contrast, there is normally more sweat involved when musicians play in turns and neither the recording nor the listening environments are great—as in most bedroom studios—since the focus natively shifts from the whole to its parts. This summons the risk that while recording, each instrument is captured so it sounds as good as it can in isolation; but during mixdown, the individual instruments may not combine well. For example, the kick and bass might sound epic when each is played in isolation but combined they would mask one another. Thinning the bass might work better in mix context. Mixing often involves altering recordings to fit into the mix—no matter how well instruments sound in isolation.

Mixing the whole versus mixing its parts

The whole versus part topic is equally important during mixdown. Do individual elements constitute the mix, or does the mix consist of individual elements? Those who believe in the former might give more attention to how the individual elements sound, but those who subscribe to the latter care about how the sound of individual elements contributes to the overall mix. It is worth remembering that the mix—as a whole—is the final product. This is not to say that the sound of individual elements is not important, but the overall mix takes priority.

> *In the world of mixing, it is common for instruments to "take one for the team."*

A few examples would be appropriate here. It is rather routine to apply a high-pass filter to vocals in order to remove muddiness and increase its definition. This type of treatment can make the vocals sound utterly unnatural when soloed. However, this unnatural sound often works extremely well in mix context. As another example, vocals can be compressed while soloed, but the compression can only be perfected when the rest of the mix is playing as well—level variations might become more noticeable with the mix as a reference. Overheads compression should also be evaluated against the general dynamics and intensity of the mix.

The question even goes into the realm of psychoacoustics—our brain can separate one sound from a group of sounds. So, for example, while equalizing a kick, we can isolate it from the rest of the mix in our head. However, we can just as well listen to the whole mix while equalizing a kick, and by doing so improve the likelihood of the kick sounding better in mix context. This might seem a bit abstract and unnatural—while we manipulate something, we seek to clearly hear the effect. The temptation to focus on the manipulated element always exists, but there is a benefit in listening to how the manipulation affects the mix as a whole.

> *It can be beneficial to employ a mix-perspective rather than an instrument-perspective, even when treating individual instruments.*

Home recordings

Opposite audio folklore stands mounting evidence that affordable equipment can deliver great results. Skill and experience are worth more than a U87. Equally, a couple of quality closed-back headphones could mitigate the need to convert the garret into a control room, let alone costing much less.

The real problem of home studios is their acoustics.

Isolation is one problem. It was extremely challenging to record any vocals during rush hour when I lived a stone's throw away from the London Overground lines. London has unique buildings and unique building materials. Some walls are made of a physics-defying material that transmits more sound than Nori. Recording becomes spectacularly problematic when on the other side of your wall is the bedroom of a young couple with a newborn. But where there's a will there's a way.

A usually bigger problem for home recordings is reverb, part of the "room sound." A typical domestic room would involve a fair amount of reverb, but more often than not, not the sort we want to capture. Acoustic treatment is never a bad thing, and it doesn't have to be professional-grade—anything that absorbs sound will reduce the amount of reverb. This includes carpets, canvases on the wall, or towels covering wardrobe doors. Microphones can also be fenced using dedicated isolation shields or even a duvet. But proper acoustic treatment can be expensive, hard to fit, and there are always diminishing returns the lower the frequencies.

Yet, generally speaking, a problematic room limits our microphone choices to those that capture less room, like most dynamic microphones. If the plot succeeds and very little ambiance is captured, this ambiance often has to be re-created during mixdown using reverbs. The good news is that even when an expensive studio is at their disposal, some producers employ a wet/dry recording approach, where many instruments are recorded with some microphones capturing the room reverb (wet) while others capturing as little of it as possible (dry). This gives more control and options during mixdown. The bad news is that for drum kits at home, it is nearly always a case of you are either unlucky or unlucky.

Another aspect of the "room sound" has to do with domestic rooms being typically small and boxy. This means that microphone placement can greatly influence the frequency content of the recordings, with the lows and low-mids being most affected (this has to do with room modes, which are discussed in chapter 8). Microphone placement, in this context, does not mean where the microphone is relative to the instrument, but rather where the microphone is in the room. It might pay to experiment with various microphone locations, with the middle of each dimension being the most likely location to avoid.

Garbage in; garbage out?

The quality of the recorded material has an enormous influence on the mix. A famous saying is "garbage in; garbage out." Flawed recordings can be rectified to a certain extent during mixing, but there are limitations. Good recordings leave the final mix quality to the talent of the mixing engineer and allow for greater creative opportunities.

Nevertheless, experienced mixing engineers can testify to how drastically the process of mixing can improve poor recordings, and how even low-budget recordings can be turned into an impressive mix. Much of this is thanks to the time, talent, and passion of the mixing engineer, and sometimes involves a few mixing "cheats," such as the triggering of drum samples and re-amping of guitars.

> *Garbage in; garbage out. Still, a lot can be improved during mixdown.*

Life has it that sometimes the quality of the raw tracks is poor to an extent that makes them impossible to work with. For instance, distorted guitar recordings that involve such a degree of comb filtering, that it makes more sense to re-record them than to fix the recording. But then, what you can't fix you can hide, trash, or break even more. For example, if an electric guitar was recorded with a horrible pedal monophonic reverb that just does not fit into the mix, maybe over-equalizing it and ducking it with relation to a delayed version of the drums will yield such an unusual effect that nobody will

notice the mono reverb anymore. Clearly there is a limit to how many instruments can receive such dramatic treatment so sometimes re-recording is inevitable.

In other cases, it is the arrangement that is to blame, as when the recorded tracks involve a limited frequency content or the absence of a rhythmical backbone. Again, there is nothing to stop the mixing engineer from adding new sounds, or even re-recording instruments; nothing, apart from time availability and ability.

What you can't fix you can hide, trash, or break even more.

Editing

Generally, on projects that are not purely sequenced, editing is the final stage before mixing. Editing is subdivided into two types:

- **Selective**—primarily concerned with choosing the right takes, and the practice of *comping*—combining multiple takes into a composite master take.
- **Corrective**—done to repair a bad performance.

Anyone who has ever engineered or produced in a studio knows that session and professional musicians are a true asset. But as technology is moving forward, enabling more sophisticated performance corrections, mediocre performance is becoming more acceptable—why should we spend money on vocal tuition and studio time when a plugin can make the singer sound in tune? Drum correction is also a common practice. On big projects, a dedicated editor might work with the producer to do this job. Unfortunately, though, sometimes it is the mixing engineer who needs to do such a job.

A lot of corrective editing can be done mechanically. Most drums can be quantized to metronomic precision, and vocals can be made perfectly in tune. Although many pop albums feature such extreme edits, many advocate a more humanized approach that calls for little more than an acceptable performance (perhaps ironically, sequenced music is often humanized to give it feel and swing). Some argue that over-correcting is against genuine musical values. It is also worth remembering that corrective editing typically involves some quality penalty. In addition, audio engineers are much more sensitive to subtle details than most listeners. To give an example, the chorus vocals on Beyoncé's "Crazy in Love" are late and offbeat, but many listeners don't notice it.

Sequenced music

The production process of sequenced music (Figure 4.3) is very different in nature to that of recorded music. In a way, it is a mixture of songwriting, arranging, and mixing—producing for short. This affects mixing in two principal ways. First, today's DAWs, on which most sequenced music is produced, make it easy to mix as you go. The mix is an integral part of the project file, unlike a console mix that is stored separately from the multitrack. Second, producers commonly select samples or new sounds while the mix is playing along; unconsciously, they choose sounds based on how well they fit into the existing mix. A specific bass preset might be dismissed if it lacks definition in the mix, and a lead synth might be chosen based on the reverb that it

Production ⟶ **Mixing** ⟶ Mastering

Figure 4.3 Common production chain for sequenced music.

brings with it. Some harmonies and melodies might be transposed so they blend better into the mix. The overall outcome of this is that sequenced music arrives at the mixing stage partly mixed.

As natural as this practice may seem, it is not without its issues. To begin with, synthesizer manufacturers and sample-library publishers often add reverb (or delay) to presets in order to make them sound bigger. These reverbs could be permanently imprinted onto the multitrack submission and have restricted depth, stereo images, and frequency spectrums that might not integrate well with the mix. Generally speaking, dry synthesized sounds and mono samples offer more possibilities during mixdown. In addition, producers sometimes get attached to a specific mixing treatment they have applied, such as the limiting of a snare drum, and print these to the submitted multitrack. Very often, the processing is done using inferior plugins, in a relatively short time, and with very little attention to how the processing affects the overall mix. Flat dynamics due to over-compression or ear-piercing highs are just two issues that might have to be rectified during a separate mixing stage.

Sequenced music often arrives at the mixing stage partly mixed—which could be more of a hindrance and less of a help.

Where to start

Preparations

Projects submitted to mixing engineers are usually accompanied by documentation—session notes, track sheets, edit notes, etc.—these are worth inspecting. Clients might have ideas, guidelines, or requirements regarding the mix, which are often discussed at this stage. In addition, there are various technical tasks that might need to be accomplished; these will be discussed shortly.

Auditioning and the rough mix

Unless we were involved in earlier production stages and are fluent with the raw tracks, we must listen to what is about to be mixed. Auditioning the raw tracks allows us to learn the musical piece; capture its mood and emotional features; and identify important elements, moments, or problems we will need to rectify. We study our ingredients before we start cooking.

Often, a rough mix (or a monitor mix) is provided with the raw tracks. It can teach much about the song and inspire a particular mixing direction. Even when a rough mix is submitted, creating our own can be extremely useful—in doing so, we familiarize ourselves with the arrangement, structure, quality of recording, and, maybe most importantly, how the musical qualities can be conveyed. Our own rough mix is a

noncommittal chance to learn much of what we are dealing with, what has to be done, and how; it can help us to start formulating a mixing plan or vision.

Rough mixes, especially our own, are beneficial.

One issue to be aware of is that it is not unheard of for the artist and/or mixing engineer to unwittingly become so accustomed with the rough mix that they use it as an archetype for the final mix. This unconscious adoption of the rough mix is only natural since it often provides the first opportunity to hear how different elements transform and begin to sound like the real thing. However, this could be dangerous territory since rough mixes, by nature, are just that—rough. They are done with little attention to small details, and sometimes involve random try-outs that make little technical or artistic sense. Yet, once accustomed to them, we find it hard without them—a point worth remembering.

Sometimes the rough mix is a throwaway; other times it serves as a skeleton for the actual mix. Such a workflow is similar to the two distinct stages employed by many artists when drawing. First, a pencil is used to quickly sketch the artwork outline, with any inaccuracies quickly corrected by "drawing over" with a bolder shade. Once satisfied, the artist then moves into a much slower phase concerned with precision and detail. This method provides some guidance as to how we may want to approach rough mixes—a quick, carefree go at getting the foundations in place. Under strict time-constraints, people tend to focus on the most valuable things, and noncommittal settings liberate us from various phycological gremlins that can get in creativity's way. Such an approach to rough-mixing is well worth trying out, as it can improve both our skills and the mix in question.

The plan

Just bringing faders up and doing whatever seems right is arguably not very effective—a bit like playing a football match without tactics. Every mix is different, and different pieces of music require different approaches. Once we are familiar with the raw material, a plan of action is developed, either in the mind or on paper. Such a plan can help, even when the mixing engineer recorded or produced the musical pieces—it resets the mind from any sonic prejudgments and creates a fresh start to the mixing process. Below is an example of the beginning of a rough plan before mixing commences. As you can see, this plan includes various ideas, from panning positions to the actual equipment to be used; there are virtually no limits to what such a plan might include:

> I am going to start by mixing the drumbeat at the climax section. In this production the kick should be loud and in-your-face; I'll be compressing it using the Distressor, accenting its attack, then adding some sub-bass. The pads will be panned roughly halfway to the extremes and should sit behind everything else. I should also try and distort the lead to tuck in some aggression. I would like it to sound roughly like "Stronger" by Kanye.

It might be hard to get the feel of the mix in the very early stages. Understandably, it is impossible to write a plan that includes each and every task. Moreover, such a detailed plan can limit creativity and chance. Therefore, instead of one big plan, it can be easier to work using *small plans*—whenever a small plan is finished, a new evaluation of the

mix takes place and a new plan is established. Here is a real-life example of a partial task list from a late mixing stage:

- Kick sounds flat
- Snare too far during the chorus—replace with triggers (chorus only)
- Stereo imbalance for the strings—amend panning
- Solo section: violin reverb is not impressive enough
- Haas guitar still not defined
- Automate snare reverbs

Not every mixing engineer approaches mixing with a detailed plan; some do it spontaneously and in their heads. Yet, there is always some methodology, or some procedure being followed.

> *Mixing is rarely a case of "whatever seems right next."*
> *Have a plan, and make sure you identify what's important.*

Technical vs. creative

The mixing process involves both technical and creative tasks. Technical tasks are generally those that do not actually affect sounds, or those that relate to technical problems with the raw material. They usually require little sonic expertise. Here are a few examples:

- **Neutralizing (resetting) the desk**—analog desks should be neutralized at the end of each session, but this is not always the case. Line gains and aux sends are the usual suspects that can later cause trouble. Line gains can lead to an unbalanced stereo image or unwanted distortion. Aux sends can result in unwanted signals being sent to effects units.
- **Housekeeping**—projects might require some additional care, sometimes simply for our own convenience. For example, file renaming, removing unused files, consolidating edits, etc.
- **Track layout**—organizing the appearance order of the tracks so they are convenient to work with. For example, putting all the background vocals or drum tracks next to each other, or having the sub-oscillator next to the kick. Sometimes the tracks are organized in the order in which they will be mixed; sometimes the most important tracks are placed in the center of a large console. Different mixing engineers have different layout preferences to which they usually adhere. This enables faster navigation around the mixer, and increased accessibility.
- **Phase check**—recorded material can suffer from various phase issues (these are described in detail in Chapter 12). Phase problems can have subtle to profound effects on sounds, and it is therefore important to deal with them at the beginning of the mixing process.
- **Control/audio grouping**—setting any logical group of instruments as control or audio groups. Common groups are drums, guitars, vocals, etc.
- **Editing**—any editing or performance correction that might be required.
- **Cleaning up**—many recordings require the cleaning up of unwanted sounds, such as the buzz from guitar amplifiers. Cleaning up can also include the removal of extraneous sounds such as count-ins, musicians' talk, coughs, etc. These unwanted

sounds are often filtered either by gating, a strip-silence process, or region trimming. (Some of these tasks might be performed at the end of the mixing stage.)

- **Restoration**—unfortunately, raw tracks can incorporate noise, hiss, buzz, or clicks. These are more common in low-budget recordings, but can also appear due to degradation issues. It is important to note that clicks might be visible but not audible (inaudible clicks can become audible if being processed by devices such as enhancers or if being played through different D/A converters). Some restoration treatment, de-noising for example, might be applied to solve these problems.

Creative tasks are essentially those that allow us to craft the mix. These might include:

- Using a gate to shape the timbre of a floor tom.
- Tweaking a reverb preset to sweeten a saxophone.
- Equalizing vocals in order to give them more presence.

While mixing, we have a **flow of thoughts and actions**. The **creative process** could do without distractions. Any technical task can break the creative flow. If, while equalizing a double bass, you find that it is offbeat on the third chorus, you might be tempted to context-switch and fix the performance straight away (after all, a bad performance can be really disturbing). By the time the editing is done, you might have switched off from the equalizing process or the creative process altogether. It can take some time to get back into the creative zone. It is therefore beneficial to go through all the technical tasks first, which clears the path for a distraction-free creative process.

At any rate, it pays to develop the discipline to defer technical tasks whilst "in the zone."

Technical tasks can interrupt the creative flow.
Better complete them first (or last).

Which instrument?

Different engineers have a different order in which they mix the different tracks. There are lots of differences in this business. Some engineers are not committed to one order or another—each production might be mixed in the order they see as most suitable. Here is a summary of some common approaches, and their possible advantages and disadvantages:

- **The serial approach**—starting with very few tracks, we listen to them in isolation and mix them first, then gradually more and more tracks are added and mixed. This divide-and-conquer approach enables us to focus well on individual elements (or stems). The danger is that as more tracks are introduced, there is less space in the mix.
- **Rhythm, harmony, melody**—we start by mixing the rhythm tracks in isolation (drums, beat, and bass), then other harmonic instruments (rhythm guitars, pads, and keyboards), and finally the melodic tracks (vocals and solo instruments). This method often follows what might have been the overdubbing order. It can feel a bit odd to work on drums and vocals without any harmonic backing. But arguably, from a mixing point of view, it makes little sense mixing in an organ before the lead vocal.
- **In order of importance**—tracks are brought up and mixed in order of importance. So, for instance, a hip-hop mix might start with the beat, then lead vocals, then

additional vocals, and then all the other tracks. The advantage here is that important tracks are mixed at early stages when there is still space in the mix and so they can be made bigger. The least important tracks are mixed last into a crowded mix, but there is less of a consequence in making them smaller.

- **Parallel approach**—this involves bringing all the faders up (of key tracks), setting a rough level balance, rough panning, and then mixing individual instruments in whatever order one desires. The advantage with such an approach is that nothing is being mixed in isolation. It can work well with small arrangements but can be problematic if many tracks are involved—it can be very hard to focus on individual elements or even make sense of the overall mix at its initial stages. As an analogy, it would be like playing football with eight balls on the pitch.

There are endless variations to each of these approaches. Some, for example, start by mixing the drums (the rhythmical spine), then progress to the vocals (most important), and then craft the rest of the mix around these two important elements. Another approach, which is more likely to be taken with more spatial mixes, involves mixing the front instruments first, and then adding the instruments panned to the back of the mix.

There are also different approaches to **drum mixing**. Here are a few things to consider:

- **Overheads**—the overheads are a form of reference for all the other drums. For example, the panning position of the snare might be dictated by its position on the overheads. Changes made to the overheads might affect other drums, so there is some advantage in mixing them first. Nonetheless, many engineers prefer to start from the kick, then progress to the snare, and only then perhaps mix the overheads. Sometimes the overheads are the last drum track to be mixed. Another factor here is how loud the overheads will end up being—in some mixes there is not much of them, so they are less of a concern for other drums.
- **Kick**—being the most predominant rhythm element in most productions, the kick is often mixed before any other individual drum and sometimes even before the overheads. Following the kick, the bass might be mixed and only then other drums.
- **Snare**—as the second most important rhythm element in most productions, the snare is often mixed after the kick.
- **Toms**—as they are typically only played occasionally, they are often the least important contributors to the overall sound of the drums (and yet their effect in the mix can be vital). Some leave the toms to be mixed at a very late stage.
- **Cymbals**—the hi-hats, ride, crashes, or any other cymbals might have a sufficient presence on the overheads. Often in such cases, the cymbals are used to support the overheads, or only mixed at specific sections of the song for interest's sake. Sometimes these tracks are not mixed at all.

Which section?

With rare exceptions, the process of mixing involves working *separately on the various sections*. Each section is likely to involve different mixing challenges and a slightly different arrangement (choruses are commonly denser than verses). And so, mixing engineers usually loop one section, mix it, then move on to the next section and mix it based on the existing mix. The question is: Which section should be first? There are two approaches here:

- **Chronologically**—starting from the top and slowly advancing to succeeding sections (intro, verse, chorus). It seems logical to work this way since this is the order in which music is played and recorded. However, while we might mix the verse perfectly—creating a rich and balanced mix—there will be very little place in the mix for new instruments introduced during the paramount chorus.
- **In order of importance**—the most important section of the song is mixed first, followed by the less important sections. For a recorded production, this section is usually the chorus; for some electronic productions, it will be the climax. Very often, the most important sections are also the busiest ones; therefore, mixing them first can be beneficial.

Which treatment?

A common guideline for treatment order is shown in Figure 4.4. With the exception of faders (which need to be up for sound to be heard), there is no reason to adhere to this. Later in the book, we will encounter a few techniques that involve a different order. However, there is some logic in this worth covering.

Faders → Pan-pots → Processors

Modulation → Delay → Reverb

Automation
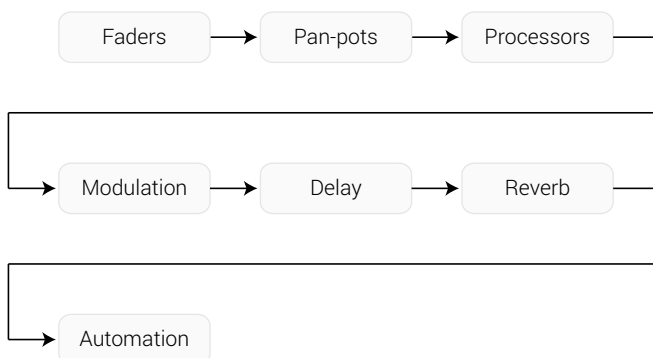
Figure 4.4 The "standard" treatment order guideline.

If we skip panning and mix in mono, we lose both the width and depth dimensions. Yet, since masking is more pronounced in mono, sometimes engineers choose to resolve masking by equalizing while listening in mono. This might be done before panning, but more frequently done using mono summing. Since processing replaces the original sound, it comes before any effects that add to the sound (modulation, delay, or reverb). The assumption is that we would like to have the processed signal sent to the effects, rather than what might be a problematic unprocessed signal. Similarly, it is usually desirable to have a modulated sound delayed, rather than having the delays modulated. Finally, since reverb is generally a natural effect, we normally like to have it untreated. In some cases, automation is the last major stage in a mix.

There is also the **dry-wet approach**, in which all the tracks are first mixed using a dry treatment only (faders, pan-pots, and processors), and only then a wet treatment is applied (modulation, delay, and reverb). This way, the existing sounds are all dealt with before adding new sounds to the mix. It also leaves any depth or time manipulations (reverbs and delays) for a later stage, which can simplify the mixing process for some. However, some claim that it is very hard to get the real feel and the direction of the mix without depth or ambiance.

The iterative approach

Back in the days of two-, four-, and eight-track tape recorders, mixing was an integral part of the recording process. For example, engineers used to record drums onto six tracks, then mix and bounce them to two tracks and use the previous six tracks for additional overdubs. The only thing that limited the number of tracks to be recorded was the accumulative noise added in each bounce. Back then, engineers had to commit their mix, time and again, throughout the recording process—once bouncing took place and new material overwrote the previous tracks, there was no way to revert to the original drum tracks. Such a process required enormous forward planning from the production team—they had to mix while considering something that hadn't been recorded yet. Diligence, imagination, and experience were key.

Today's technology offers hundreds of tracks. Even when submix bouncing is needed (due to channel shortage on a desk or processing shortage on a DAW), the original tracks can be reloaded at later times and a new submix can be created. This means that practically everything can be altered at any stage of the mixing process, and nothing has to be committed before the final mix is printed.

Flexibility to undo mixing decisions at any point in the mixing process is a great asset since **mixing is a correlative process**. The existing mix should normally be retouched to accommodate newly introduced tracks. For example, no matter how good the drums sound when mixed in isolation, the introduction of distorted guitars into the mix might require additional drum treatment (the kick might lose its attack, the cymbals might lose definition, and so forth). Plus, treatment in one area might require subsequent treatment elsewhere. For instance, when brightening the vocal by boosting the highs, high frequencies may linger on the reverb tail in an unwanted way, so the damping control on the reverb might need to be adjusted. The equalization might also make the vocal seem louder in the mix and so the fader might need to be adjusted. If the vocal is first equalized and then compressed, the compression settings might need to be altered as well.

Since mixing is such a correlative process, it can benefit from an **iterative coarse-to-fine approach** (Figure 4.5). This means starting with a coarse treatment on which we spend less time, then, as the mix progresses continually, refining previous mixing decisions. Most of the attention is given to the late mixing stages, where the subtlest changes are made. There is little justification in trying to get everything right before these late stages—what sounds fine at one point might not be so later.

Start with coarse and finish with fine.

Milestones

The mixing process can have many milestones. On the macro level, there are a few that are potentially key (Figure 4.6). The first involves bringing the mix into an adequate state. Once this milestone is reached, the mix is expected to be free of any issues—whether they existed on the raw tracks or were introduced during the actual process of mixing. Such problems can range from basic issues such as relative levels (e.g., solo guitar too quiet) to more advanced concepts such as untidy ambiance.
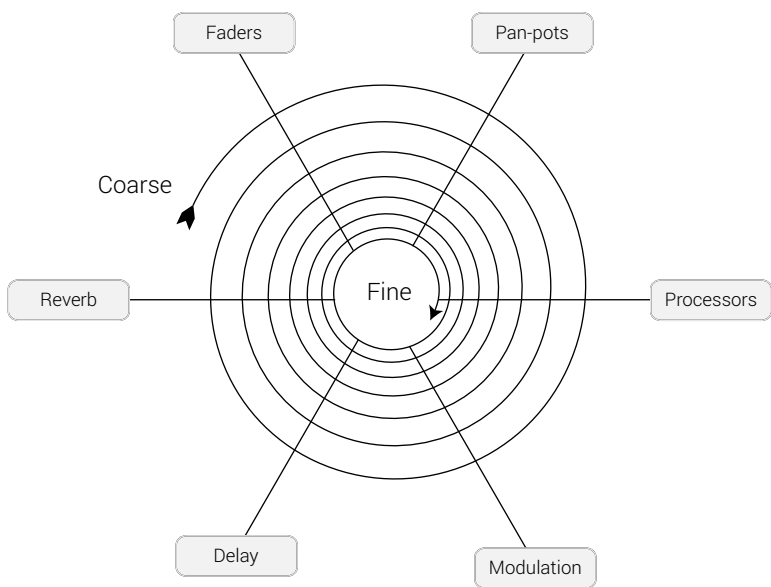
Figure 4.5 The iterative coarse-to-fine mixing approach.
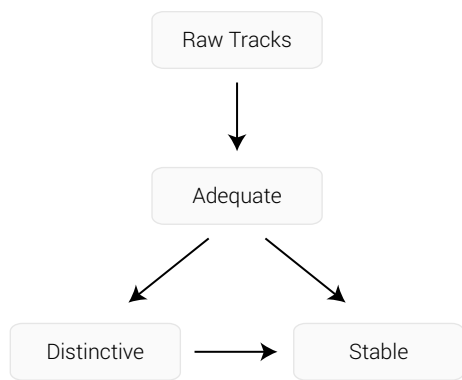
Figure 4.6 Possible milestones in the mixing process.

Nevertheless, a problem-free mix is not necessarily a good one. The next milestone often involves making the mix distinctive or appealing. The definition of such qualities are abstract and these vary between one mix and another, but the general objective is to make the mix notable—memorable—whether by means of interest, power, feel, or any other sonic property.

The final milestone, which often only applies to the inexperienced engineer, is stabilizing the mix.

Finalizing and stabilizing the mix

All the decisions we make while mixing are based on the evaluation of what we hear at the listening position; otherwise known as the *sweet spot*. But mixes can sound radically different once one leaves this sweet spot. Here are some prime examples:

- Different **speakers** reproduce mixes differently—each speaker model (sometimes in combination with its amplifier) has its own frequency response, which also affects perceived loudness. While listening on a different set of speakers, the vocals might not sound as loud and the hi-hats might sound harsh. Different monitor positions also affect the overall sound of the mix, mainly with relation to aspects such as stereo image and depth.
- Mixes sound different in different **rooms**—each room has its own sonic stamp that affects the music played within it. Frequency response, which is affected by both room modes and comb filtering, is the most notable factor, but the room response (reverb) also plays a part.
- Mixes sound different in different **points in a room**—both room modes and comb filtering alter our frequency perception as we move around the room. A mix will sound different when we stand next to the wall and when we are in the middle of the room. (This issue is discussed in greater detail in Chapter 8.)
- Mixes sound different at different **playback levels**—as per our discussion on equal-loudness curves.

Many people are familiar with the situation where their mix sounds great in their home studio but translates badly when played in their car. It would be impossible to have a mix sounding the same in different listening environments for the reasons explained above. This is worth repeating: mixes *do* sound very different in different places. What we can do is ensure that our mixes are problem-free in other listening environments. This is much the idea behind stabilizing the mix.

The reason that it is usually the novice that has to go through this practice is that the veteran mixing engineer knows her mixing environment so well that she can predict how well the mix will sound in other places. Also, in comparison to a professional studio, the sweet spot in a home studio is actually quite sour. Acoustic problems in home studios can be profound, while in a professional studio these are rectified by expensive construction and acoustic design, carried out by experts.

The process of stabilizing the mix involves listening to the mix in different locations, different points within those locations, and at varying levels. Based on what we hear, we can finalize the mix by fine-tuning some of its aspects. Normally, this involves subtle level and equalization adjustments. Here is how and where mixes should be listened to during the stabilizing phase:

- **Quiet levels**—the quieter the level, the less prominent the room and any deficiencies it exhibits. Listening quietly also alters the perceived loudness of the lows and highs, and it is always a good sign if the relative level balance in the mix hardly changes as the mix is played at different levels. Listening at loud levels is also an option, but can be misleading as mixes usually sound better when played louder.
- **At the room opening**—also known as "outside the door." Many find this listening position highly useful. When listening at its opening, the room becomes the sound

source, and instead of listening to highly directional sound coming from the (relatively) small speakers we listen to reflections that come from many points in the room. This reduces many acoustic problems caused by the room itself. The monophonic nature of this position can also be an advantage.

- **Car stereo**—a car is a small, relatively dry listening environment and many find it useful to audit their mixes there. In the same way as the room-opening position, a car provides a very different acoustic environment that can reveal problems.
- **Headphones**—with more and more people listening to music on mobile devices via headphone, comes the growing need to check mixes using headphones. With the use of headphones, we will sacrifice some aspects of stereo image and depth but not have to worry about room modes, acoustic comb filtering, and left/right phase issues. We can regard listening via headphones as listening to the two channels in isolation, and it is known that headphones can reveal noises, clicks, or other types of problems that may not be as noticeable via speakers. The main issue with headphones is that you need a good pair, and that the stereo panorama changes quite drastically compared to a playback via speaker (having said that, nowadays listeners are more likely to listen to the mix on headphones than on a pair of well-positioned speakers).
- **Specific points in the room**—many home studios present poor acoustics, and very often there are issues at the sweet spot. Moving out of this spot can result in an extended bass response and many other level or frequency revelations. Also, the level balance of the mix changes as we move off-axis and out of the monitors' close range. Beyond a certain point (the critical distance), the room's combined reflections become louder than the direct sound—most end listeners will listen to mixes under such conditions, so it is important to hear what happens away from the sweet spot.

Evaluating the mix using any of these methods can be misleading if caution is not exercised. While listening in a specific point in a room, it might seem that the bass disappears and subsequently we might be tempted to boost it in the mix. But there is a chance that it is that point in the room causing the bass deficiency, while the mix itself is fine. How, then, can we trust what we hear in any of these places? The answer is that we should focus on specific places and learn the characteristics of each. Usually, rather than going to random points, people evaluate mixes in the exact same point of the room every time. In addition, playing a familiar reference track to see how it sounds in each of these places is a great way to learn the sonic nature of each location.

While listening to the mix in different places, it could be wise to **make a note of** what issues you notice and which require further attention, then go back to the mixing board. In some cases, comments cancel one another out. For example, if while using head-phones it seems that the cymbals are too harsh but when listening at the room opening they sound dull, perhaps nothing should be adjusted.

As a final piece of advice, listening in different locations is not reserved for final mix stabilization—it can also be beneficial at other stages of the mixing process, especially when decisions seem hard to make.

Mix edits

Very often, especially in the commercial industry, more than one version of a mix is required; these versions are referred to as "edits." Primarily, this practice is concerned

with creating a mix that will conform to the destination playback system. Some of these edits can be produced during the mastering stages. But mixing engineers have more control over individual tracks and, hence, the final result. Common edit candidates are:

- **Album version**—the mix to be included on the album.
- **Radio edit**—various factors should be considered for a mix that is intended for broadcast on the radio. First, since mixes played on the radio are heavily compressed and limited before transmission, mixing engineers sometimes check their mixes through heavy compression or limiting to see how the mix will translate under such conditions. Second, since it is fair to assume that radio listeners are likely to listen in noisy environments, vocals are commonly pushed up in radio edits. Third, longer songs are less likely to be played on commercial radio and therefore long album versions are commonly shortened using edits and fades. Fourth, since most radio systems have limited ability to reproduce low frequencies, very often these are filtered or attenuated to some extent. Finally, some lyrical content might require censorship, which is usually done by mutes or 1 kHz tone.
- **Club and LP versions**—it is assumed that both of these will be pressed on a vinyl, which requires centered bass content, and minimum phase differences between the left and right channels, especially at low frequencies. As opposed to radio, club sound systems are expected to provide an extended low-frequency response, so mixing engineers must use full-range monitors to make sure that the low end is properly mixed. Most clubs have a limiter on the signal chain as well.
- **Vocals-up/vocals-down**—the levels of the vocals in mixes are critical. Very often, two mixes are bounced, with the vocals varying by around 1–2 dB between the two. The A&R, producer, and artist usually pick their favorite version. If appropriate, it is also possible to record additional variations, such as drums-up/drums-down.

As well as these common edits, additional edits or stems might be required; for example, an instrumental mix, a cappella mix, video mix, TV (instrumental and backing vocals), no solo, and so on.

Mastering

Why master?

Why do we need mastering engineers? Primarily because anyone about to release a track to the public would like the last person to work on that track be a highly skilled and experienced engineer that works in a studio with optimal acoustics. Mastering is partly about ensuring quality and increasing quality and we need a specialized person in a specialized studio to do that.

Mastering engineers have more roles than meet the eye. When assembling an album, they remove extraneous sounds, arrange the tracks in the most compelling order, apply smooth fades and natural pauses, and balance both the frequency spectrum and the level of the various tracks so the album sounds like a coherent piece rather than a collection of unrelated songs.

Once a master is completed and approved by the client, they produce a high-quality copy that is fit for digital streaming, and, in some cases, one that complies with the requirements of manufacturing plants.

Perhaps their most important role is to bring the sonic aspects of an album to their best and most appealing state. If the mixes are good, mastering engineers can make polished diamonds out of gold.

The equipment used in a professional mastering studio usually cost more than its equivalent in mixing facilities, and the listening environment is optimized to rectify possible problems, mostly acoustic. It is common, for example, to find mastering studios with nothing that could cause comb filtering (including a desk) between the full-range monitors and the listening position. Theoretically, mastering engineers might have to amend the mixes very little if the mixing engineer did the job right. But rarely do mixing engineers have the environment or tools to achieve the critical quality that mastering engineers can.

It should be clear why mastering is so significant—once the finished master leaves the mastering studio, any problems will be heard by many and will potentially damage commercial success, sales, and—most importantly—the joy of listening.

> *Simply put, mastering is an art and science reserved for the experts.*

Some mixing engineers submit mixes that have some mastering-like treatment, mostly compression. In subtle amounts, and if done correctly, it can improve the overall sound of the mix. But common sense has it that whatever a mixing engineer can do on the stereo mix a mastering engineer can do better. Why would you try to fix your company car if your company will pay a professional mechanic to do it for you?

Mastering delivery

Historically, mixes used to be submitted for mastering on ½″ analog tapes. Later, DATs became popular, then CD-ROMs, DVD-ROMs, external hard drives, USB sticks, and nowadays many just upload files to the cloud. Although uncommon today, CD-DAs (audio CDs) are prone to errors more than other types of media, and should not be used. The actual submission is often accompanied by some documentation (track names, etc.).

When working on a stereo mix, one of the greatest challenges in mastering is that each instance of processing affects all the elements of the song. For instance, correcting sibilant vocals can reduce the snare's clarity when sibilance occurs. The high-fidelity tools at the mastering engineer's disposal can fix many issues, but the more the mix needs correction, the more quality is compromised. Since nowadays it is possible to find multitrack applications in a mastering studio, it is not uncommon to submit mixes in *stems* (submixes of logical track groups, or even just a single track). If there are any problems in the mix, the mastering engineer might find it easier to process the individual stems rather than the mix as a whole. Common stems are vocals, rhythm, leads, and the residue mix (which consists of everything but what is already included in other stems). In all cases, a full stereo mix should be submitted as well, and should be identical to the mix of all the other stems when their faders are at unity gain (0 dB).

Sometimes the client or mastering engineer asks for changes to mixes after these have been completed. In a large studio with an analog console, recalling mixes can be time-consuming. Saving a mix in stems can be beneficial in such situations—instead of recalling the whole mix (console and outboard gear), we can only recall the mix of the stem that requires alterations, while playing all the other stems untouched.

There are a few additional practices worth considering when submitting mixes to mastering:

- **Do not fade**—leave any fades at the beginning or end of each track for the mastering engineer. He or she can do more-musical fades once the order of the tracks has been determined, and can use any noises at the beginning or end of the track for noise reduction. Make sure that you leave the full reverb tail at the end of the track or any other instrument decay.
- **Use WAV files**—these are supported on both Mac and PC platforms, and aren't compressed.
- **Keep the original digital audio quality**—do not perform any sample-rate or bit-depth conversions. These are likely to degrade the quality of the audio, and have no advantage from a mastering point of view—many mastering engineers will convert the mixes to analog before processing, and will use high-quality converters to capture the analog signal and then convert it back to the appropriate digital format.
- **Disable any master bus processing**—sometimes the master bus in a mixing session has EQs, compressors, limiters, enhancers, or other processors. These can make the mix more appealing during mixing sessions, but most likely better redone during mastering. If terribly attached to these, one can always submit two versions—with and without master bus processing.
- **No clips**—ensure the master bus audio doesn't clip.

The loudness war

Louder is perceived as better. We have already established that. In a world where everyone wants to be better, everyone wants to be louder. And once the meme has spread globally, it eventually got tagged as a "war." It is nothing short of ludicrous that too often the main ask from mastering engineers is to make it "as loud as possible"; and it is equally sad that many online mastering services are judged by the loudness of the masters they return, not by the overall quality of the masters. But where there's a human bias, there will be those to exploit it. The catch is that crazy-loud masters are likely to sound worse to end listeners than anything sensibly loud or even tracks that are not loudness-boosted at all.

Radio stations were one of the first to exploit the commercial prospect of loudness. If you switch from 99FM to 100FM and the latter sounds better, you are less likely to tune out. If you then switch to 101FM and it doesn't sound as good, you are more likely to switch back. Add advertising to the mix and loudness equals money. FM radio stations had limiters installed already because overshooting levels meant a penalty for overshooting the station's allocated frequency range. Making your station louder simply meant pushing the volume against the limiter. But any station could do that, and the more you push the more distortion, and things can turn nasty quick. So radio stations started incorporating ever more sophisticated chains of loudness boosting in order to prevail. Nobody wanted to be left behind, not even classical stations, which explains

why the same classical recording sounds better on the radio than when played from a CD at home. If fact, critical listeners might notice that songs mastered before the loudness war could sound better on the radio than very loud masters from after the loudness war; old masters get louder by the station's loudness processing and will be as loud as any other track, whereas loud masters might distort. Regardless of how loud a radio station is on the scale, radio stations need to ensure that all songs are roughly at the same level—a practice known as *normalization*; so regardless of how loud your master is, you should expect it to sound roughly as loud as any master on professional radio stations.

All the of above happened in an analog world, with loudness being increased while flowing through analog units. Things are different in today's digital world. Today, we can take any audio file and meassure how loud it is. A unit called *LUFS* (Loudness Unit Full Scale) captures exactly that, and it involves a somewhat complex algorithm that measures rather faithfully *perceived* loudness. Streaming services—like Spotify or YouTube—also need to ensure songs are at the same level, so each stream all songs at a fixed LUFS level (which is typically in the range of −16 to −9 LUFS). If a master is louder than that fixed level, it will be attenuated; if it is quieter, it will be boosted (barring few services that only attenuate). So here again, no matter how loud your master is, it will end up at the same level as any other song when streamed.

Up until very recently, nearly all music streaming over the Internet was done using lossy compressions algorithms (like AAC or MP3). This allows streaming using lower bandwidth, which reduces costs. Specifically costly is digital transmission through airwaves, so digital radio stations often sacrifice quality by using even lower bitrates than streaming. The lower the bandwidth, the more quality degradation, specifically with louder masters.

So the tables have turned, with crazy-loud masters sounding worse than their sensible counterparts. So does that mean we shouldn't boost the loudness of our mixes during mastering? Not at all. Not only a loudness boost can enhance sonic qualities, but quiet masters can end up being quieter on some playback chains. But it should never be a case of making a master loud for the sake of winning a competition—nowadays, the winners of the loudness war are nearly always losers.

Any loudness boost during mastering should stop when the quality of the master starts degrading. It is probably wise to keep masters within −16 to −9 LUFS, but even this depends on many other factors (the nature of the track and its most likely playback mediums). There are various LUFS meters out there, and various Codec previewers out there (for example, by iZotope or even Apple) and masters should be checked with these.

Finally, LUFS allows us to measure the perceived loudness of audio like never before, which facilitates a fair comparison when it comes to loudness. If we produce two masters—one with −14 LUFS and another with −8 LUFS—we should expect the latter (the louder) to sound better. But we can now compare them more objectively by attenuating the latter by 6 dB and thus both will be played at −14 LUFS; only then can we compare the quality of both.

DIY mastering

Mastering at home can be as challenging as eating soup with a fork in a pitch-dark room—hard to do without the right tools and setting. But that doesn't mean there aren't ways to produce better masters at home.

Compensating for the lack of

As discussed above, a highly acoustically calibrated room and good full-range monitors are paramount for mastering, but amiss in pretty much every domestic room. We just need to work harder; the same advice given for stabilizing the mix applies to mastering as well. First, and admittedly controversial, a good (which in this case means relatively expensive) set of headphones could be better than a room with flawed acoustics. Then, we can check the masters in other settings.

The bottom end, however, is not that easy to get right as it is generally impossible to treat what we cannot hear. Few have access to a studio with full-range monitors. The bottom end is also where terrible accidents can happen. Some genres make use of sub-bass and during mixing it will be boosted so to actually hear its effect; as much of that range is rolled off by the near-field speakers, there is some likelihood it will be boosted disproportionally. Then play the mix via full-range monitors and the sub-bass overpowers the whole mix. Although not appropriate to all genres (club music, for example), one option is to apply a high-pass filter, starting with the lowest cut-off frequency available and slowly moving it up until we can hear its effect. This will filter any low frequencies we cannot hear. A better safe-than-sorry approach. The good news is that this sort of treatment is sometimes applied by mastering engineers anyhow—since most consumer speakers and headphones cannot reproduce very low frequencies, a mix with these rolled-off could sound better on a wider range of devices.

It is intuitive to assume that the superior equipment in mastering suites translates into better masters. When it comes to the acoustics and monitors, this is true, but less so with processors. In a particular academic experiment done in the noughties, a world-renowned mastering engineer was asked to master a few tracks using the in-house equipment. Following that, the same engineer was asked to produce identical masters using Logic plugins only. Then a group of university staff and students were asked in a "blind" test to tell which master is the in-house one and which is the Logic one. The results were quite surprising—a statistical analysis showed that everyone was pretty much guessing (across all tracks). This research was not rigorous enough for its results to be considered valid, or make it to publication; yet one would expect people would be able to easily tell.

Instead of focusing on the cost of the processors found in mastering suites, it is worth focusing on their nature. Mastering suites do not have the same quantity of processors as mixing studios, with much of the mastering chain using selected few hardware units. Also, many of these units do not have continuous controls, but rather step ones (which makes recall easier). Mastering engineers have *less* options, which provides them with what is known as *enabling constraints*. It is easier to learn when there is less to learn, and easier to make choices when there are fewer options. This advice will be repeated in this book, but the less-is-more maxim applies to mastering as well.

Mastering tools

Equalizers together with compressors and limiters are key to mastering. Unlike in mixing, the settings on these are typically "gentle," with equalizers applying a mere 0.5 dB of boost or attenuation in some cases. Since each treatment affects all the instruments in the mix, the overall effect of any processor is more drastic.

Equalizers are used to fix excess or deficiencies on certain ranges, or just to give the track its fitting frequency balance—hip-hop would get more bass, retro-rock less so.

Compressors are normally configured rather conservatively. The general aim is to push the mix against the compressor to increase its loudness and to level-up the various instruments. The dynamics of the instruments at some point will start taking a hit, with what's appropriate depending on the genre and the artist. The upper limit of this treatment is normally when pumping occurs. Since a mix involves the full frequency spectrum, side-chain equalization can help attain better results, but at the cost of extra complexity.

Loudness maximizer plugins[1] are an alternative to compressors and limiters when wishing to boost the loudness of the mix. These typically require less expertise than compressors and could produce fair results. Most maximizers also apply brick-wall limiting to prevent clipping.

Many analog limiters are somewhat compressors with less controls, so these can be used to achieve the same task as compressors. Digital limiters are true limiters, and generally a finer control would be provided by a compressor. Practically every mastering chain has a true limiter as the last processor (barring dither) to contain clips.

Gates and duckers are rare in mastering.

Reverbs are not common in mastering, although a property-configured quality reverb at a subtle amount can add a bit of life to a mix, specifically an overly dry one. Delays and other modulation tools are uncommon.

Stereo width manipulation is unlikely to be part of DIY mastering as any changes can easily be applied to the mix itself.

Analog devices are flawed but in a way most find appealing. The subtle distortion they add tends to improve masters, and as no stereo unit can process both channels in an identical manner, minor differences between the left and right channels tend to open up the sound slightly. None of this happens with digital processors (unless intentionally added). We could add a subtle amount of distortion using distortion plugins (probably best done in a parallel configuration). On some DAWs, it is possible to process the left and right channel separately, so dialing ever so slightly different settings on the left and right channels can mimic the analog artifact. But possibly a better option is to use a dedicated enhancer or a good emulation plugin.

A fundamental chain for DIY mastering could be:

• A compressor
• An equalizer

- An enhancer
- A loudness maximizer
- A brick-wall limiter (should be included in a loudness maximizer)
- Dither (could be included in a loudness maximizer or a limiter)

The compression and equalization are where most skill is needed and most time would normally be spent.

Why split the process?

A possible question at this point would be: If we are both mixing and mastering at home in the same room, why split the process? Why not just master while mixing by putting some plugins on the master bus? That's not a crime, but caution should be exercised, specifically with compressors and limiters, as they can make mixes sound better than they are and reduce our ability to discern important details. When a hard-working compressor is used throughout the mixing process, one should not be surprised to discover that the level balance of the mix is all over the shop when these are disabled. Another thing worth remembering is that despite mastering our own tracks, there is some chance that in the future we may want (or be asked) to have them mastered by professionals; there is also a chance we will need to compile individual tracks onto an EP or an album. In these situations, having non-mastered mixes is handy.

Furthermore, the focus of mixing and mastering is quite different. During mixing, we care about how the various instruments interact; during mastering, we care about how the mix sounds as a whole. This seems to contradict the previously mentioned advice that we should try and listen to the mix as a whole while mixing; this stands, but not in the same way as in mastering. For example, it is possible to get recordings that are bass heavy; when we mix those, we try for get all the instruments to gel together, but we care less about the fine details of the overall mix frequency spectrum—this is something done during mastering. For sure, if the mix is clearly bass heavy, we may roll off some low frequency from this or that instrument, but we are not paying the same amount of attention (nor have the studio do so in most cases) as we would when mastering. The same goes for the overall loudness of the mix and many other sonic treatments that are usually done during mastering. One way to think about it is that mixing is about the *relative*—how instruments combine compared to each other, whereas mastering is about the *absolute*—how the mix compares to many similar tracks out there.

A/B nearly like a pro

When mastering engineers assemble an album, they need to ensure that sonic qualities like levels and frequency content are consistent between songs (within artistic reason). One way to achieve this is to first master the "anchor" track, possibly the lead single. Once done, the master is printed and serves as a reference to all other tracks. When working on another track, mastering engineers will occasionally toggle to the anchor track for comparison and amend any sonic differences. Practically speaking, the currently mastered track will be loaded to a stereo track on a DAW, the output of which goes through a mastering chain. The (printed) anchor track will be on another track, the output of which is unprocessed.

A similar setup can be used for DIY mastering. We can pick a commercially mastered track we aspire to mimic and use it as the anchor track. Then, while mastering our own track, repeatedly refer to it and try to match the overall sonic qualities. Matching the loudness of a commercial track may not be attainable, so sometimes the reference track will have to be attenuated.

Note

1 Naming isn't easy and often marketing takes precedence; the world is full of bad names, which sometimes we get so used to that we just don't notice. The so-called loudness maximizers do not maximize loudness because there isn't quite a maximum level you can bring your track onto. Strictly speaking, if you *really* maximize your audio, you will end up with 1-bit audio sound, which is not something anyone would like to listen to. A better name may be Smart Loudness Boosters, or SLBs in short. If they aren't smart, then they are limiters, yet some limiters are smartish.

5 Related issues

How long does it take?

A B-side of a commercial release can be mixed in as little as three hours. An album's lead single might take six days. Generally, for a standard release, it would be normal to spend a day or two per song; within this period, an engineer will often nail the core of the mix (the critical sonic decisions and overall feel) within a few hours and anything after this could be of a slightly more technical nature. Albums are very often mixed within about three weeks; Nirvana's *Nevermind* took about 14 days to mix. The complexity of the track is an obvious factor—a 24-track production that has a simple structure, involving a five-piece rock band, should take less time than a 96-track production with many varying sections. The recording quality also plays a part—poor recordings mean that the mixing engineer spends more time fixing the recordings than creatively crafting the mix.

There is a difference between mixing in a commercial studio and mixing at home. Studio engineers are normally restricted by deadlines. But there is very little to restrict the home producer from mixing one production for a whole month. This can be both an advantage and a disadvantage. Many professionals will testify that they hardly ever properly finish a mix—they just bring it to its best state within a given time constraint. There is always something to improve, always a chance to add more excitement or impact. So, for the novice, not working to strict deadlines can be a good thing. But mixing a song for anything more than a week can easily mean it is being over-mixed. Many engineers say that part of the wisdom of mixing is to know when to jump from the carousel, or it will keep spinning forever.

If Kate finishes a mix in two days and Axl finishes a mix in a month, Kate will mix 182 songs a year while Axl will mix only 12. Kate is more likely to be the better mixing engineer after a year. What's more, when working for many days on the same production, it is hard to remain excited or maintain the same state of mind throughout.

However, it is important to remember that, unlike the veteran, the novice is walking a long and intense learning path—while the veteran is fluent with their equipment and mixing environment, the novice has much knowledge and experience to acquire. The veteran may choose a compressor and compress the lead vocal in a matter of seconds (having done so daily for years). The novice might have to spend time going through different compressors, trying them in isolation and then with the mix, perhaps while also reading the manual.

But there is some pleasure to be gained from being a novice. Discovering mixing and constantly coming up with better mixes can give great satisfaction. For the veteran, mixing is an occupation—a daily job that might involve projects she is not fond of, busy schedules, and annoying clients. But the beauty of mixing is that every project differs from the next, and even after 30 years of mixing there is always something to learn.

Some people question how professional engineers can spend eight hours a day mixing. Most of them do not—the music industry is so demanding that it is very common for a mixing engineer to be working in excess of 12 or sometimes even 18 hours a day. This might seem inconceivable for beginners, who after just a few hours can find themselves "unable to hear anything anymore." Ear-fatigue is common and, even when monitoring levels are kept low, the brain can simply get tired after a long period of attentive listening. But with time, our brains learn to handle longer mixing periods.

Deadlocks

The evaluation block

This is probably the most common and frustrating deadlock for the novice. It involves listening to the mix and sensing that something is wrong, but not being able to specify what. Reference tracks can help in these situations—they give us the opportunity to compare aspects of our mix with an established work. This comparison can reveal what is wrong with our mix, or at least point us in the right direction.

The circular deadlock

As we mix, we define and execute various tasks—from small tasks such as equalizing a shaker to big tasks such as solving masking issues. The problem is that in the process, we naturally tend to remember our most recent actions. While in the moment we might not pay any attention to a specific compression applied two days before, we might be tempted to reconsider an equalization applied an hour ago. Thus, it is possible to enter a frustrating deadlock in which recent actions are repeatedly evaluated—instead of pushing the mix forward, the mix goes around in circles. A possible way out of this situation is a short break. We might also want to listen to the mix as whole, and re-establish a plan based on what really needs to be done.

Be mindful of reassessing recent mixing decisions.

Breaks

Sometimes the process of mixing can be so enthralling that time flies by and we haven't even felt the need to take a break. But, as many recording engineers will testify, long sessions without a break put a strain on us, and can make us less effective as time goes by. Mixing is a brain- and ear-demanding process that requires attentive listening,

which can be hard to maintain for long periods. Breaks help us to forget our recent mixing actions and therefore provide an opportunity to move the mix forward. After a break, we can also bring down the monitoring level if it was brought up earlier—doing so after a break seems somewhat easier than bringing the level down in the middle of the process, when we are used to hearing it at the louder level.

Probably the most important break, for the novice anyway, is the **critical break**—a day or two without listening to the mix after completing it. Having such a long break can neutralize our brain and make us forget individual treatments we have applied. After a long break, we are likely to listen to the mix with less sonic prejudice.

The danger with solos

Solo buttons let us listen in isolation to specific tracks. A useful function all right, but one that can be easily misused, especially by the novice, who might spend too much time mixing isolated tracks. We have already established that a mix is a composite and that it can be beneficial to adapt a mix-perspective approach. Using solos results in the opposite since it promotes element-perspective approach. No doubt soloing a track makes changes easier to discern, but there is also no doubt that it can lead to a treatment that is out of mix context.

One of the problems in using solos is that we lose a reference to the mix. An example already given would be compressing vocals when soloed—they might sound balanced in isolation, but not appear so with the rest of the mix. When soloed, nothing acts as a reference to the loudness of the vocals. Soloing additional tracks, say the acoustic guitar, can give such reference. But, as other instruments might affect the perceived loudness of the vocals, only compressing them with the rest of the mix ensures solid results. On the same basis, panning or depth positioning is pointless unless done in respect to the rest of the mix.

There are situations where using solos is very sensible, such as when trying to filter out the buzz from a guitar track. Also, sometimes it is hard to focus on a very specific component of the sound with the whole mix playing along; for example, the resonant frequency of a snare. It may be useful to look for such frequency when the snare is soloed, however, it would be pointless in keeping the solo once the frequency has been found.

Nonetheless, solos should be used whenever we manipulate sounds with relation to themselves; for example, while applying EQ automation on vocals in order to balance out frequency changes caused by the proximity effect.

Solos should be used with caution since they promote
an element-perspective approach.

Mono listening

Mixes are often checked in mono. We should first consider where music plays in mono:

- **Television**—some televisions are still monophonic.
- **Radio**—standard AM radio broadcast is monophonic, while stereo FM receivers switch to mono when the received signal is weak. Like with televisions, cheap models might only involve a single speaker.
- **Large venues**—it is very hard to maintain a faithful stereo image to large audiences in a big venue. Moreover, in many of these venues, there is very little point in distributing a stereo signal, and such a setup would be more expensive. Therefore, in places such as malls, sport stadiums, and supermarkets, the stereo signal is summed to mono prior to distribution.

The term *mono-compatible mix* is used in relation to mixes that translate well when summed to mono. Unless mixes are clearly intended to be played in mono (which is fairly rare), the aim is not to create a mix that will present no issues when summed to mono, but to **minimize the side effects** that this summing might bring about. In practice, there are endless examples of great mixes that present some issues when summed to mono—the loss of high frequencies on Nirvana's "Smells Like Teen Spirit" is just one example. Mixing engineers may sacrifice mono-compatibility for the benefit of a powerful stereo effect, but they will also make sure that the issues caused by mono summing are minimized. A good example of this practice relates to the delay times used for the Haas trick (see Chapter 12)—different time settings result in different comb-filtering effects when the stereo material is summed to mono.

> *We usually listen in mono to minimize mono problems,*
> *not to eliminate them completely.*

Many studios have a single speaker installed for mono listening. Having the sound only coming out of one speaker is considered *true mono*, as opposed to the slightly blurred *phantom mono* produced by two speakers. As already mentioned, some engineers find it easier to resolve masking when listening in mono. However, it is worth knowing that when we sum to mono, the balance of the mix changes. Why exactly this happens is explained in Chapter 14. For now, it is sufficient to say that the center remains at the same level while the extremes drop by 3 dB.

Track 5.1: Hero Stereo
The "Hero" production in stereo.

Track 5.2: Hero Mono
The same production in mono. Most affected are the electric guitars, which can hardly be heard here.

Mono listening also helps in **evaluating the stereo aspects of our mix**—by switching between mono and stereo, it can be easier to determine the authenticity of panning, various stereo effects, and the overall impact of the stereo panorama. A slight change between stereo and mono might suggest that more stereo processing should take place.

Many analog and digital desks offer a mono switch. Very often in the software domain, a specific plugin with such a functionality has to be inserted on the master bus (Figure 5.1).

Figure 5.1 Logic's Gain plugin enables the summing of a stereo signal to mono.

Housekeeping

There have been cases in the past where studio assistants lost their jobs for failing to label a tape or for not completing a recall sheet. Some very simple (sometimes boring) tasks can save enormous amounts of time and effort afterward. Housekeeping is one of them. The word "many" is encountered frequently during the process of music creation—we deal with many clients, many reels, many tracks, many files, many takes, many versions, and so forth. These many things are much easier to manage, work with, identify, trace, or recall if proper housekeeping is exercised.

Take labeling, for example: a file labeled *track 12_tk9* suggests nothing about its content. A file labeled *Kick_tk3* does. But worse is a file labeled *Kick_tk3* that actually contains a vocal. Going through a project that consists of 200 files (which is not unusual) and encountering meaningless or incorrect file names is every engineer's nightmare.
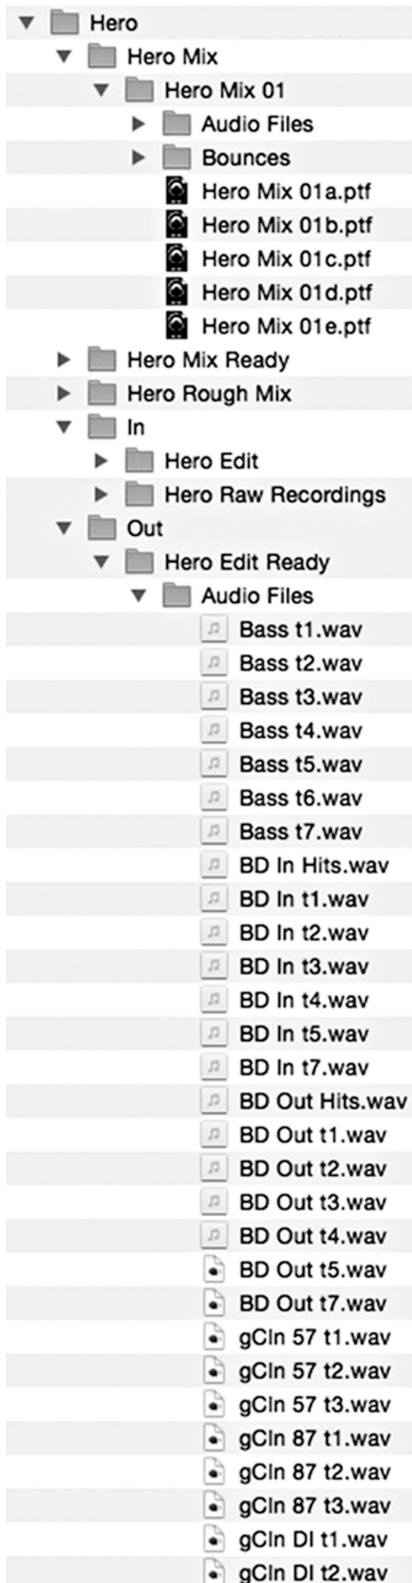
▼ 📁 Hero
  ▼ 📁 Hero Mix
    ▼ 📁 Hero Mix 01
      ▶ 📁 Audio Files
      ▶ 📁 Bounces
        📄 Hero Mix 01a.ptf
        📄 Hero Mix 01b.ptf
        📄 Hero Mix 01c.ptf
        📄 Hero Mix 01d.ptf
        📄 Hero Mix 01e.ptf
  ▶ 📁 Hero Mix Ready
  ▶ 📁 Hero Rough Mix
  ▼ 📁 In
    ▶ 📁 Hero Edit
    ▶ 📁 Hero Raw Recordings
  ▼ 📁 Out
    ▼ 📁 Hero Edit Ready
      ▼ 📁 Audio Files
        Bass t1.wav
        Bass t2.wav
        Bass t3.wav
        Bass t4.wav
        Bass t5.wav
        Bass t6.wav
        Bass t7.wav
        BD In Hits.wav
        BD In t1.wav
        BD In t2.wav
        BD In t3.wav
        BD In t4.wav
        BD In t5.wav
        BD In t7.wav
        BD Out Hits.wav
        BD Out t1.wav
        BD Out t2.wav
        BD Out t3.wav
        BD Out t4.wav
        BD Out t5.wav
        BD Out t7.wav
        gCln 57 t1.wav
        gCln 57 t2.wav
        gCln 57 t3.wav
        gCln 87 t1.wav
        gCln 87 t2.wav
        gCln 87 t3.wav
        gCln DI t1.wav
        gCln DI t2.wav

Figure 5.2 This screenshot shows a partial folder hierarchy for the "Hero" project presented in this book. The "Hero Mix" folder contains the actual mix project. There might be more than one mix per song, so there is a subfolder called "Hero Mix 01." The various session files are snapshots of the mix-in-progress. The "Hero Mix Ready" contains a project from which a new mix can start; it is a modified version of the edited version, involving the desired track order, groups, and so forth. The "Hero Rough Mix" folder contains the rough mix. The "In" folder is an archive for incoming material; it includes the raw recordings and the editor's submission. The "Out" folder is an archive folder for material submitted to other entities; in this case, the "Hero Edit Ready" folder contains the version submitted to editing, which is a modified version of the raw recordings (file names, comments to editor, etc.). Notice the straightforward audio file names, with perhaps the exception of "gCln 57 t1," which stands for "guitar clean *SM57* take one."

6 Mixing and the brain

Psychology and cognitive sciences have been subject to some exciting breakthroughs in the last few decades, perhaps most notable of which is the emergence and widespread acceptance of the theory covered in this chapter. *Dual process theory* has been immensely fertile, providing coherent explanations to both individual and social phenomena, and settling millennia-long debates. And all of this is done while carrying the hallmark of a great theory—simplicity.

Being useful to many disciplines, the theory can also be used to explain some important aspects related to mixing. This chapter is centered around psychology and is largely conceptual, so readers more interested in hands-on mixing can safely skip it and read it later.

Dual process theory

We generally have a strong sense that we're in control, that we know what is going on in our heads, why we act how we act, and why we choose what we choose. Yet, we all sometimes shed a tear without quite planning to, feel down for no apparent reason, shout at our computer screens (fully aware they can't hear us), or buy things we don't really need.

As far as historians can trace it, human beings have suspected that part of our mind—part of us—is not within our control. So who or what is sharing the driver's seat with our conscious mind, and to what extent does it affect our thoughts and behavior?

From Plato to Freud, this question has kept philosophers and thinkers up at night. It shaped early Hindu thinking, fashioned into the biblical Satan, and is the subject matter of countless fictional works, such as *The Strange Case of Dr Jekyll and Mr Hyde*.

In the latter half of the twentieth century, scientists set sail to unravel the mystery, and by the new millennium a body of research was all pointing the finger at one thing: Our brain. More specifically, a multitude of theories combined with mounting empirical evidence have made a convincing case that processes within our brain can be classified into two unique categories:

- **Unconscious**—automatic, fast, effortless, involuntary, emotional, and intuitive processes.
- **Conscious**—controlled, slow, effortful, voluntary, rational, and analytical processes.

> *In a research paper from 2000, Keith Stanovich and Richard West summarized a multitude of theories dealing with these two types of processes. In order not to commit to this adjective or another, they intentionally used the generic term system 1 to describe unconscious processes, and system 2 to describe conscious ones. The scientific world adapted these terms, which become popular among the general public with Daniel Kahneman's publication of his best-selling book* Thinking, Fast and Slow *in 2011. But with consciousness being better defined and understood today than in 2000, like many recent publications, this book uses the terms* unconscious *and* conscious *together with* processes *or* mind*.*

Let's get to know these two types of processes better.

Two process types

What do you see in Figure 6.1?

Figure 6.1 Something.

You saw a face. If asked how do you know it is a face, your answer may be along the lines of: "It has two eyes, a nose, and a mouth so it must be a face." But such an answer is really retrospective reasoning—a theory you've made up. Can you tell whether it is a male or a female? A baby or an adult? How do you know? You just do.

In fact, as soon as your eyes landed on this page, you immediately knew it was a face of a young adult female. I'd even venture to suggest you classified her as pretty. There was no *conscious* thinking process involved in all of this; rather, it was all worked out by your unconscious mind. The recognition process was fast, effortless, and involuntary—you simply couldn't help seeing it is a face.

Now allow yourself 30 seconds to try to fill the missing number in the following sequence:

0, 3, 8, ?, 24

In this case, your unconscious mind is clueless. You'd have to work this out, probably while attempting various strategies, before these words may echo in your head: "Take 1–5, square each number, and subtract one." A process that is conscious, effortful, and slow.

From conscious to unconscious

Many unconscious processes are innate. The most inner part of our brain—called the reptilian brain, because we share it with reptiles—is the host for the processes that keep us alive, such as those regulating our heartbeat and body temperature. Other processes we were programmed to acquire involuntary very early in life; these typically give us an evolutionary advantage, such as face recognition, which is crucial for the formation of functional human societies.

But other kinds of unconscious processes have nothing to do with survival or evolution—we acquire them later in life, and do so through conscious, effortful practice. A popular example is that of driving a car. Initially we expand great conscious effort to execute the sequence of actions driving involves. But with practice, driving becomes an automatic affair—we no longer think of when to press the clutch, shift the gear, or release the break. The same transition of conscious to unconscious processes can be observed with many abilities, such as tracing a line, reading, playing an instrument, or… mixing.

When I was teaching undergraduates compression, I used to start the set of lectures with an experiment. After providing some background on levels, dynamics, and compression, I played the class three masters in random order. One had no compression, the other moderate compression, and another was *Death Magnetic* style—"squashed to death," so to speak. Students were asked to rank the three masters based on how compressed they are. For one with trained ears, the differences were glaring. But those with no experience with compressors or musical dynamics could not tell the difference. This is much like how people with no experience in photography cannot identify an overexposed photograph. When a similar experiment was conducted after students had spent 20 hours with compressors, virtually everyone in the class got the ranking right, and did so easily. When asked, "How could you tell?" the reply was always the same: "You just hear it."

The same can be said of many aspects of mixing, such as hearing the room on a recording, identifying resonating frequencies, getting gating right, or the operation of various tools. What at the start requires concentration and conscious effort with practice becomes easy and automatic.

The power of the unconscious

The human brain is the most powerful and efficient processor in the universe. Despite being the size of a large grapefruit and consuming the energy of an energy-saving lightbulb, compared with the most powerful supercomputer in the world—Frontier—our brain is roughly 1,800 times faster. So how exactly is such vast processing power distributed between the conscious and unconscious minds?

It turns out that our unconsciousness is by far the more powerful one of the two. There are a few key reasons for this. First, reasoning—a conscious process—is

normally carried out in verbal form, so there's a limit to how fast the voice in our head can "speak." Second, it was shown that conscious *thinking* (as opposed to *awareness*) is the result of back-and-forth communication with the unconscious, with each cycle taking up to 200 ms; if our thought process requires feedback from the world around us, another 300 ms are lost (300 ms is roughly the time it takes information to reach from our senses to our consciousness). Third, our brain can only have one conscious process at a time, whereas unconscious processes are many and run in parallel.

Our unconscious brain is capable of truly phenomenal feats. How long does it take to solve a Rubik's Cube? For people with no experience, who will do so purely consciously, it would take roughly 5 minutes given the instructions are right in front of them. The current world record is less than 5 seconds. How is it possible? Rubik's masters rehearse the process until the rules settle in the unconscious mind like the rules of grammar (which native speakers follow without any conscious thought), the algorithms become as native as words, the moves as letters, and the whole execution is done using muscle memory—completely bypassing consciousness. All of this is the workings of the unconscious, with the conscious mind nearly exclusively dedicated to "look ahead"—figuring out the next algorithm.

Similarly, chess masters don't quite consciously analyze future move paths—even for five turns, the amount of possible board combinations is upwards of a billion. They do consciously employ strategies and heuristics, but most of the game is based on an unconscious "sense" of the board. Adding up 15 three-digit numbers when they are flashed in less then 2 seconds? No problem for those in Japan who practice mental abacus—participants can hardly recall any of the numbers, but still get the answer right.

All of these and many other wowing skills would not be possible had it not been for our unconsciousness, and—in the cases above—countless hours of practice. Rubik's Cube, chess, and mental abacus masters practice for years, often more than 4 hours a day.

The literature on mixing is full of recipes—do this, use that. But recipes are followed by the conscious mind. There is no doubt conscious thinking is an important part of mixing—everyone has their procedures, techniques, and reasons. But unlike the instructions on how to solve the Rubik's Cube, no one can give you a set of instructions ensuring a great mix. Much of mixing is the workings of the unconsciousness—the real powerhouse of the brain. To harness its power is relatively easy and fun—mix.

Exercising the unconscious

Conscious processes turn unconscious with practice. But not all practice is equally fertile—fairly early on, driving a car doesn't make one a better driver anymore, unless one is a racing driver. Once a process turns automatic, it hardly improves with repetition, if at all. Although our unconscious mind changes even without interaction with the conscious mind, experiments have shown that the conscious process—a massive and powerful activation event in our brain—has far greater capacity shaping unconsciousness. This links to a famous theory called *levels of processing*; it asserts that the more conscious effort we expend, the higher the likelihood of understanding and

retention. This is why solving a problem yourself will teach you more than seeing how it is solved, why you are more likely to remember material by summarizing it before an exam rather than just reading it, and why reading about mixing has only a slight chance of improving your skills without applying hands-on what you've read about.

So improving our unconscious requires a conscious effort. But what level of effort is optimal? The answer to this originated through the eyes of a 10-year-old boy fleeing the deadliest war in history.

Warm guns

It is 1944; Bulgaria; World War II is raging. On the last outgoing train to cross the Danube River sat a 10-year-old boy called Mihaly Csikszentmihaly (pronounced MEE-high CHEECH-sent-mee-highee). The boy had a nagging thought—he observed that, given their circumstances, some people who are supposed to be unhappy are happy, and some who are supposed to be happy aren't. This observation led that boy to later conceive and dedicate his life to a new science—the science of happiness. After years of research, he presented his findings and theories in a book called *Flow: The Psychology of Optimal Experience.*

Neuroscience teaches us that whenever we conquer a challenge, a feel-good chemical is released in our brain—a simple reward mechanism. Brain damage, dysfunction, or drug intervention can impair this mechanism, resulting in apathy and misery. Generally speaking, the harder the challenge, the higher the reward.

Is doing the dishes enjoyable? Most people will see it as a burden. But flow theory asserts that it can be enjoyable so long as some *achievable* challenge is set. Can you do it faster? Can you better coordinate your moves? Use less water?

For nearly a century now, some of the best scholars of our time have attempted the formulation of a single consistent theory that explains our physical universe, to no avail. Had you managed this, you'd be ranked in history books higher than Einstein. A challenge all right. But can you do it? Is it achievable?

In accord with the reward mechanism described above, flow theory asserts that hard challenges are more rewarding than easy ones. But flow theory also accounts for our skills. A challenge too easy for our skills will result in apathy. A challenge too hard, will result in anxiety rather than reward. Csikszentmihaly's flow model is shown in Figure 6.2. Of particular interest is the state of flow—a state reported by many artists, professionals, and scientists—that involves complete immersion in the task at hand, a lost sense of time, mental elevation, and a great sense of fulfillment. Csikszentmihaly has shown that such a state is the result of a high challenge combined with high skill, and that it can be maintained over time if the challenge involves small and manageable increments in difficulty.

Cognitive science offers much support to the flow theory—posing a high challenge to the conscious mind but ensuring our skills (our unconsciousness) can manage it, thus gaining high reward. As the consciousness and unconsciousness dance together in the process, skills form and strengthen in our unconsciousness, allowing even harder challenges to be conquered. And the cycle continues.
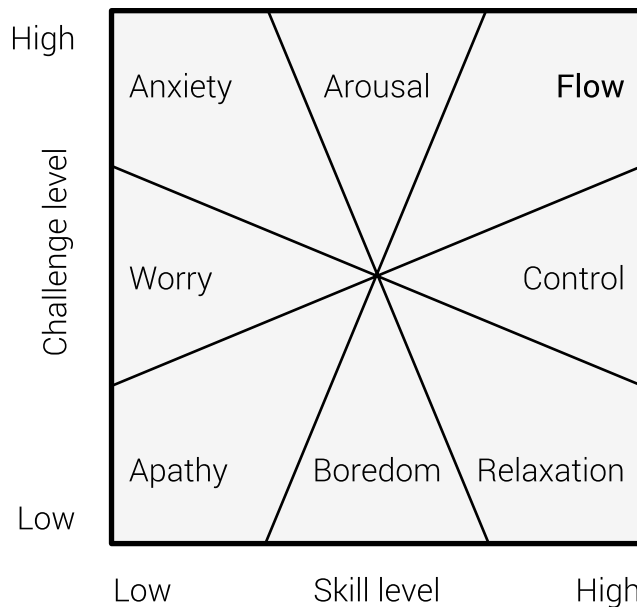
Figure 6.2 The flow model.

There are two main takeaways from this with relation to mixing. First, if you have only recently entered the world of mixing, do not delve into despair if your mixes do not pan out to be world class at first. World-class mixes require many unconscious skills that take time to acquire. It is hard enough to understand how compressors work, an understanding that is largely the outcome of conscious effort. But the fine ability to hear dynamic variations, their characteristics, and how each control on a compressor might affect them is largely an unconscious skill.

Second, structure your practice so it involves manageable challenges that gradually increase in difficulty. Well-recorded and arranged productions are easier to mix. The same is often the case with sparse arrangements. It is easier to start with material of limited diversity and utilize a few, rather than many, tools. There is no need to browse all the 256 reverb presets—short-term memory limits would only allow you to properly survey around five at a time. You may just as well limit your choice to three presets: the default hall, room, and plate. Only once you are proficient with the characteristics and utilization of this triplet, venture to increase your reverb arsenal.

Intuition

I sometimes had to mix in front of students, sessions that involved thinking aloud. For sure, part of the process involved a conscious analytical approach that can easily be reported—"bass varies in levels and has resonating notes, it needs compression, possibly a multi-band one" or "a spring reverb would probably work with this vocal, giving it a bit more vintage feel." But more often than not, especially when twiddling knobs, I have found myself incapable of explaining why I'm doing what I'm doing, even to myself. It is as if some internal agent is guiding you, telling you what to do. Psychologists call such an agent *intuition*, and needless to say it's completely unconscious.

When one says, "I feel she dislikes me," one means that his or her impression is based on intuition, rather than reason. Interpersonal communication is hugely intricate and relies on many subliminal cues, so no wonder we often base our assessment on intuition. But as we all know—and as scientists have proven—intuitions are sometimes wrong. Combined with the fact that Western philosophy is obsessed with reason, intuition has somewhat of a bad reputation, more so in the West. Eastern philosophy is far more forgiving to intuition; internal feelings, unexplained as they are, are not as easily dismissed. Interestingly, it is Western scientists that have shown that intuition, while terribly imperfect, is also of considerable capabilities.

The psychologist Gary Klein has studied decision making under pressure. In a famous case he studied, a commander urged his firefighters to evacuate a burning house for no apparent reason. Soon after the firefighters escaped, the floor they stood on collapsed. When asked why he made the call, the commander had no immediate explanation other than the feeling something was wrong. Only after the event, a coherent explanation was devised—a massive fire in the basement atop which the firefighters stood made the room they were in atypically hot for the fire noise within. It was the unconscious mind of that commander, trained through experience, that saved the firefighters' lives.

The same intuition, or "sixth sense" if you wish, is evident in many cases where parents *feel* their child is unwell, yet struggle to tell why. Behavior is also intricate.

A paradigm emerges here. Problems that can be solved by algorithms are fit for the conscious mind. But problems involving high complexity—too many variables that cannot be easily captured, coded, and reasoned about—are better left to the unconscious. Respectively, what computers can and cannot do.

So where is mixing in all this? Again, parts of it clearly call for conscious thinking. But if we run the same mathematical analysis that chess games are subject to, one can easily show that given the amount of possible input states (the board configuration in chess, what we hear in mixing) and possible moves (e.g., which frequency to boost and by how much), at any given point in mixing one has more options than the amount of electrons in the universe. Mixing has inherent complexity that calls for the participation of the unconscious mind.

A frequent advice in mixing is, "If it sounds good, it's good." As far as a logical argument goes, it is complete nonsense: What sounds good is subjective; within a society, people will disagree on subjective matters (otherwise these will be objective); even a single individual constantly changes so what sounds good now may not sound as good in an hour's time. But the slogan isn't really a rational argument. Instead, it captures a particular wisdom related to the unconsciousness—"trust your intuitions."

The two modes of mixing

When it comes to mixing, dual process theory has it that each of us can employ two principal modes. The first involves rational conscious thought: "I hear X, therefore I should do Y"; we'll term this mode *mixing-by-reason*. The second involves listening to your intuition, what your unconscious mind tells you: "It sounds good"; we'll call it *mixing-by-intuition*.

People who learn about mixing tend to rely more on mixing-by-reason. After all, new knowledge is acquired by conscious means and has to be recalled and put to the test through the same channel. When starting to mix, mixing intuition is hardly developed to reliably inform decisions. But the latter statement overlooks something. After all, most of us acknowledge that hi-fi speakers sound better than those of mobile phones, and we seem to exhibit a preference for singing in the bath than in an open field. Even people who don't even know what mixing is often pass judgment on the sound of a song or a gig as great or terrible. This explains the many stories about professional engineers who lack fundamental theoretical knowledge, such as what decibels are, the Q control on an equalizer, or the attack and release on a compressor. So even without mixing experience, one can, to quite an impressive extent it seems, mix-by-intuition.

But the mixing intuition each of us seems to inherently possess grows as we mix, waiting to be employed as a powerful guide, if only we choose to "go with our feelings." The danger is to overlook its powers and be forever stuck in a mix-by-reason mode. There are a few good reasons why people show a tendency to get stuck this way. One of them is that reason itself serves us with reasons as to why we do what we do; a form of security that we are doing the right thing. You don't quite know what your intuition will tell you or how to explain your actions afterwards—basing our decisions on something we have no introspection on or control over is somewhat risky.

Some of us are more rational, while others are more connected with their intuitions (or emotions). This manifests itself in many walks of our life. To give one example, when visiting a restaurant for the first time, how do you choose a dish from the menu? Do you narrow down the choices, then semi-systematically compare the remaining candidates by subscribing fors and againsts to each option? Or do you quickly spot one or two dishes that seem attractive and go with your feelings—"This looks good, I'll go with that"?

As far as mixing advice goes, be aware that we are all susceptible to a disproportionate time spent mixing-by-reason while neglecting mixing-by-intuition. The latter can be of immense power, more so with experience.

Before moving on, here is something we'll come back to later. There isn't a right or wrong answer here, just whatever first pops into your head. Complete the missing letter as you see fit:

SO_P

Thinking without thinking

Recent studies attribute even greater power to the unconscious mind. It has been theorized that our brain has two modes of operation—the *executive* mode is when our cognitive processes are controlled by our consciousness, dealing with whatever task or thought is most important at the moment. The second mode is *dreaming*—when our consciousness lets go, leaving the brain to roam in the realm of seemingly random thoughts. We all know we dream at night, but recently it was shown that we momentary toggle into daydreaming far more than we realize.

A few days of sleep deprivation and one's mental function is reduced to a level that would label one insane. The jury is still out on the particulars of why we sleep. But few contest that while dreaming, the day's experiences, with all the memories and knowledge they encompass, are integrated within our brain. In a way, a dreaming brain thinks: "I have all these new things from the last 18 hours, let's order and make long-term sense out of them now." Long-term memory is largely gained while dreaming, and part of the process involves the formulation and strengthening of links between existing and new concepts. Odd as it may seem, such process partly constitutes thinking, which in turn yields new or clearer understandings. Based on these very ideas, science finally has a fairly reasonable explanation to a previously unexplained phenomenon reported by many—Eureka moments, where novel ideas or solutions seem to pop when we *don't* think about them.

Well before these theories crystallized in psychology, thinkers and researchers in the field of creativity and problem-solving already had a concept known as the *incubation phase*—after consciously collecting related information (the preparation phase), one should let go of the problem at hand by attempting relaxation and restrain from deliberate thinking. Hopefully, such an incubation phase will lead to an illumination phase, which in turn will require a verification phase. Nowadays, the incubation phase idea has a strong cognitive backing—practically inciting a daydreaming mode to allow unconscious processes to carry out valuable work.

As paradoxical as it may sound, this suggests that in an attempt to succeed in something that requires conscious effort, one could benefit from *stopping* making such an effort periodically. When mixing, perhaps sometimes just letting go, sitting back while imagining sweet Mother Nature, and just letting the mind wander is the best thing to do. It is another way to let our powerful unconscious take its course.

Emotions

There isn't an agreement on the definition of the term *emotions*. But it is agreed that their triggering is uncontrolled—they stem from our unconsciousness.

What makes a great song or a great mix? If you ask the critics, you'll get an array of answers in reply. Critics typically have vast knowledge in their field and understand musical works within both historical and cultural contexts. They have an arsenal of analytical tools allowing them to see deep into each work. So, at least in part, their critique is based on conscious thought.

But critics are humans, and all human judgment can be heavily swayed by the unconsciousness, to a surprising extent sometimes. Countless experiments have demonstrated that. Asking someone to hold your cup of hot coffee would make you in their eyes a warmer person. In the absence of others, a mere glance on an image of a face will make individuals act more honestly; a famous study has shown that judges are far more likely to grant parole after their lunch, compared with before it. If we are already on the topic of food—shortly after being exposed to words such as *restaurant*, *food*, or *dishes*, people are more likely to complete SO_P as SOUP rather than SOAP. After all of these experiments end, and their nature and findings are revealed to the

participants, they often dismiss them: "That cup of coffee had no effect—he *was* a warm person!" But statistics tell a different story. So whether they acknowledge it or not, critics base their critique on uncontrolled emotions, to some extent at least.

The majority of the consumers that listen to the music we mix hardly respond to it by conscious means. Their response is largely, if not wholly, emotional. This, in fact, is one of the reasons we listen to music—it has this ability to provoke emotions and alter our mood in seemingly magical ways. A *mental drug*, if you wish.

If a listener's response is largely emotional, as mixing engineers shouldn't we make an effort to understand the emotional offerings of a production? And shouldn't we let our own emotions guide us in the process?

Change

Change and attention

How often do you wake up in the morning surprised by the fact you are exactly where you went to bed the night before? And while reading this book, did the thought occur to you at any point that whatever object you are reading it from hasn't disappeared? In both cases, our consciousness is aware, but not attentive (you are aware you woke up where you went to bed, but you don't spend a second thinking about it).

We sometimes feel we are living in a chaotic world, but such chaos is of immense order compared with true chaos—where you wake up, not where you went to bed, and where objects just disappear and appear at random. For life to emerge and survive, order is imperative. People who suffer from schizophrenia—a mental disorder affecting 1 percent of the population—may experience hallucinations and report disordered reality. An orderly world is not enough—one needs to perceive it as such.

Pattern recognition is a fundamental brain faculty. As we are exposed to the world, our unconscious mind constantly maintains and updates a representation of what's normal—how things *should* be. So long as things are normal, it's business as usual. But when something changes, if the change is noticeable or salient, our unconsciousness raises an alarm bell, asking our consciousness to attend.

This mechanism solves the old paradox known as the *cocktail party effect*—how can one notice what someone else says without listening? (The cocktail party effect has another definition, as we'll see later in the book.) You may stand at a party with many people chatting around you. Being engaged in a conversation yourself, you are not paying attention to anything said around you—the noise is there, but if asked what people around you were talking about you'd have no clue. But the moment someone utters your name, your attention is drawn.

The cocktail party effect is an example of a change drawing our attention. It works the opposite way as well—a lack of change and attention withdraws. Mechanical clocks make ticking sounds. If we share enough time in the same room with a mechanical clock, we stop hearing it. In fact, if you attend to the sounds in your environment right now, you may hear a noise you didn't hear a few seconds ago.

The fact that a change draws and maintains attention manifests itself in many ways in music. Nearly all productions have a structure consisting of varying sections; there are chord and key changes, and even a drum fill can be considered as an attention drawer. Without changes, people are more likely to tune out. While mixing, we can either emphasize existing changes (say, by automation), or we can add them—changing effects, adding distortions, or removing frequencies as the song progresses.

So far, we have discussed change at a micro-scale—individuals and short time periods. But change is also a huge factor on a macro-scale—when dealing with society and long periods of time.

Change is progress

I could hardly believe it was 2006 when I saw Slayer playing at Reading Festival. The group members had aged a bit all right, but how they dressed and moved, the riffs, the sound, and pretty much everything else could easily make one believe that it was 1986. Nostalgia is great. But I've found more thrill, fresh sound, and creative leads in a live set by Vitalic.

Trends come and go, artists rise and fall, musical genres pop and pop. Why? Why is music in constant flux? To begin with, society changes—few deny a casual link between World War II and the explosion of rock 'n' roll that followed; how our generation feels differs from how previous generations felt. Technology changes as well—the Moog Modular, MPC, and Ableton Live are a mere few examples of innovations that broadened our production arsenal and changed the music we listened to. But our brain plays a part as well—too long listening to the same musical ideas and we lose interest. Music would be in dire straits if 30 years on our charts were filled with "Sunday Bloody Sunday" soundalikes. We thrive on change.

If we could travel back in time to 1969, take a copy of Beyoncé's "Single Ladies" and play it to people; the likelihood is they would find it hard to digest and would probably categorize it as experimental. When a change is too big, the little that is left of the familiar may be insufficient for people to connect with the music—audiences may find the music difficult to "understand" and enjoy (in cognitive terms, such a situation can be branded as *insufficient pattern matching*).

Change is often *evolutionary* rather than *revolutionary*. But we are often misled to think the opposite, since only a few of us are experts enough to be conversant with the small steps that led to seemingly super-innovative works. Take, for example, Charles Darwin's theory of natural selection. Did he single-handedly come up with it out of the blue? At the time, Scottish philosophers led by Adam Smith came up with a new, exciting theory about how people optimize gains within society, effectively erecting a new discipline we now call *economics*. Historians argue that Darwin merely applied the same economic principles within a different context. (As imperative in science, a wealth of work is still required to tweak and prove a theory, and Darwin fully deserves the credit for the unprecedented work he has carried; it's just that the idea itself wasn't as original as many believe.) The production of Beyoncé's "Single Ladies," hailed for its creativity by many, can be seen as an evolution from a previous single she released—"Get Me Bodied"; the choreography in the video clip was inspired by a 1969 dance routine from *The Ed Sullivan Show*, and the clip shares some visual signatures with that of Claude François's 1966 "J'Attendrai"—the French version for "Reach Out (I'll Be There)."

The theory proposed by Adam Smith and his peers, now termed classical economics, *has gone relatively unchallenged for more than a century. But throughout the twentieth century, cracks have started to appear.*

*One of the key assumptions of classical economics is that people are **rational**. Based on this premise, Apple should have never succeeded. This is not to say that people who buy Apple products are irrational. Instead, Apple's reliance on emotional design, and its focus on high usability and aesthetics, should have had no effect on consumers based on classical economics theory. Apple's financial reports tell a different story.*

Within the framework of dual process theory, Daniel Kahneman and Amos Tversky have spent decades of research to show that people's behavior, judgment, and decisions are highly affected by a multitude of biases that stem from the unconscious. These biases are so strong that even knowing about them is often insufficient to override them through conscious effort. Within today's scientific community, one will be mocked for claiming that people are rational—we are simply not.

Despite being psychologists, Kahneman and Tversky would have gone on to win a Nobel Prize in Economics in 2002, had it not been for Tversky dying eight years earlier. Being alive, it was only Kahneman who was awarded the prize. Their work—offering robust validation of dual process theory—founded what is now called behavioral economics, *and it does well to explain Apple's success.*

Creativity

Creativity is defined as *the ability to produce new ideas that are of value*. Scientists are unanimous that everyone is creative—whenever you come up with a possible solution to any problem you face, you have generated a new idea that is of value. Even when coming up with a way to remember where you leave your house keys constitutes creativity. But such an idea may be original and of high value for you, but not so for mankind. To be considered creative within society, an idea needs to be new and of value for many. With fair ease, a scientist that has a new idea can find out if anyone has had such an idea before—scientific publications are searchable and incorporate systematic referencing. We don't have such privilege in the arts. Yet, years of exposure to music, especially if scoped to a particular area (say either a genre or a practice such as mixing), can give us a fair clue as to what is the norm and what will constitute a deviation from it. People sometimes fail to see how historical and cultural knowledge of their field is pivotal to their creative potential—a crazy idea for you will be a mundane one if many thought of it before you without you knowing it. There is no chance you'll get investment today for "inventing" sliced bread. The assertion that new ideas need to have a value simply comes to guard against new things that are pure nonsense. If you are to mix purely by dice throwing, you may end up with something extremely new, but unpleasantly so. Some sonic aesthetic must be maintained.

Many associate creativity with a spark of genius. But research on creativity has shown that it can be successfully pursued systematically. Analogy, randomness, opposites, or combination are all central to various creative techniques meant to aid generating new ideas. A random choice of an effect on a random mix element may work. Will beat-like

sounding drums on a country song never work? And what might it sound like if you invert the order of processors treating the vocal track? I'm sure you can come up with many more ideas.

Creativity, or the lack of it, is something that music industry people readily identify, whether they are A&R, radio editors, journalists, or the like. Aspects such as song-writing ability, musicianship, performance, arrangement, production, and *sound* are all considered. Those who go down in the history books are often the pioneers (The Beatles), not the copycats (The Monkees).

If you study the development of artists, you may identify three phases:

- **Intuition:** As we are just starting, we have little knowledge so we act primarily on intuition. When I first used compressors, I had little idea of how they worked; I just experimented and stopped when it sounded fair.
- **Education:** When one deliberately takes on some form of learning, say though reading a book or taking a course. For me, having read about compressors in *Sound On Sound* meant I could reach more musical results faster.
- **Creativity:** Once one learns the rules, it is time to break them. Venture into unfamiliar territories of new sounds and original techniques.

Some jump straight into the second phase; a few never bother with it. But it's the third phase that separates the best from the rest of us. If you are just starting to mix, your main concern is probably to get your mixes on par with what's out there. But once competent enough, remember that for music to move forward, it needs people fueling it with new ideas.

7 Mixing domains and objectives

Mixing can be seen as involving two main concerns: macromixing and micromixing. **Macromixing** is concerned with the overall mix; for example, its frequency balance. **Micromixing** is concerned with the individual treatment of each instrument; for instance, how natural the vocals sound. When we judge a mix, we evaluate both macromixing and micromixing. When macromixing, there is generally a set of areas and objectives that should be considered. These will also affect micromixing.

The process of mixing can be divided into five main domains or core aspects: **time, frequency, level, stereo, and depth**. Two of these aspects—stereo and depth—together form a higher domain: **space**. We often talk about the mix as if it exists on an imaginary sound stage, where instruments can be positioned left and right (stereo) or front and back (depth).

The term stereo *does not necessarily denote a two-channel system. Any system that is not monophonic can be regarded as stereophonic; for example, Dolby's 4.1 stereo. For convenience, however, throughout this book, the term* stereo *means a two-channel system.*

In many cases, instruments in the mix are fighting for space; for example, guitars and vocals that mask one another. Each instrument has properties in each domain. We can establish separation between competing mix elements by utilizing various mixing tools to manipulate the sonic aspects of instruments, and their presentation in each domain.

Mixing objectives

The key objective with mixing is to best convey the performance and emotional features of a musical piece so that it moves the listener. As the saying goes: There are no rules. If a mix is exciting, who cares how it was done? But there is a wealth of experience available and strategies that have been devised to help those who are yet to develop that guiding inner spark or mixing engineer's instinct, or to help grasp the technical aspects of mixing.

We can define four principal objectives that apply in most mixes: **mood, balance, definition, and interest**. When evaluating the quality of a mix, we can start by

considering how coherent and appealing each domain is and then assess how well each objective has been satisfied within that domain. While such an approach is rather technical, it provides a system that can be followed and that is already part of the mixing process anyway, even for the most creative engineers. Also, such an approach will not detract from the artistic side of mixing, which is down to the individual. First, we will discuss each of the objectives and then we will see how they can be achieved. We will also discuss the key issues and common problems that affect each domain.

Mood

The mood objective is concerned with reflecting the emotional context of the music in the mix. Of all the objectives, this one involves the most creativity and is central to the entire project. Heavy compression, aggressive equalization, dirty distortion, and a loud, punchy snare will sound incongruent in a mellow jazz number, and could destroy the song's emotional qualities. Likewise, sweet reverberant vocals, a sympathetic drum mix, and quiet guitars will only detract from the raw emotional energy of an angsty heavy metal song. Mixing engineers that specialize in one genre can find mixing a different genre like eating soup with a fork—they might try to apply their usual techniques and familiar sonic vision to a mix that has very different needs.

Balance

We are normally after balance in three domains:

- frequency balance;
- stereo image balance; and
- level balance (which consists of relative and absolute—which will be explained shortly).

An example of frequency imbalance is a shortfall of high-mids that can make a mix sound muddy, blurry, and distant. However, when it comes to the depth domain, we usually seek coherency, not balance. An in-your-face kick with an in-your-neighbor's-house snare would create a very distorted depth image for drums.

There are generally two things that we will consider trading balance for—a creative effect and interest. For example, in a few sections of The Delgados' "The Past That Suits You Best," Dave Fridmann chose to pan the drums and vocal to one channel only, creating an imbalanced stereo image but an engaging effect. This interest/imbalance trade-off usually lasts for a very brief period—rolling off some low frequencies during the break section of a dance tune is one of many examples.

Definition

Primarily, definition stands for how distinct and recognizable sounds are. Mostly, we associate definition with instruments, but we can also talk about the definition of a reverb. Not every mix element requires a high degree of definition. In fact, some mix elements are intentionally hidden; for example, a low-frequency pad that fills a missing frequency range but does not play any important musical role. A subset of definition also deals with how well each instrument is presented in relation to its timbre—can you hear the plucking on the double bass, or is it pure, low energy?

Interest

On an exceptionally cold evening late in 1806, a certain young man sat at the Theatre an der Wien in Vienna, awaiting the premier of Beethoven's *Violin Concerto in D Major*. Shortly after the concert began, a set of four dissonant D-sharps were played by the first violin section and the orchestra responded with a phrase. Shortly after, the same thing happened again, only this time it was the second violin section, accompanied by the violas, playing the same D-sharps. On the repetition, the sound appeared to shift to the center of the orchestra and had a slightly different texture. "Interesting!" thought the young man.

We have already established that our brain very quickly gets used to sounds, and without changes we can lose interest. The verse-chorus-verse structure, double-tracking, arrangement changes, and many other production techniques all result in variations that grab the attention of the listener. It is important to know that even subtle changes can give a sense of development, of something happening, even though most listeners are unconscious of it. Even when played in the background, some types of music can distract our attention—most people would find it easier studying for an exam with classical music playing in the background rather than death metal. Part of mixing is concerned with accommodating inherent interest in productions. For example, we might apply automation to adapt to a new instrument introduced during the chorus.

A mix must retain and increase the intrinsic interest in a production.

The 45 minutes of Beethoven's violin concerto mentioned above probably contain as many musical ideas as all the songs that a commercial radio station plays in an entire day, but it requires active listening and develops very slowly compared with a pop track. Radio stations usually restrict a track's playing time to around three minutes so that more songs (and advertisements) can be squeezed in. This way, the listener's attention is held and they are less likely to switch to another station.

Many demo recordings that have not been properly produced include limited ideas and variations, and result in a mundane listening experience. While mixing, we have to listen to these productions again and again, and so our boredom and frustration are tenfold. Luckily, we can introduce some interest into boring songs through mixing—we can create a dynamic mix of a static production. Even when a track has been well produced, a dynamic mix is often beneficial. Here are just a few examples of how we can increase dynamism in a recording:

- Automate levels.
- Have a certain instrument playing in particular sections only—for instance, introducing the trumpets from the second verse.
- Use different EQ settings for the same instrument and toggle them between sections.
- Apply more compression on the drum mix during the choruses.
- Distort a bass guitar during the choruses.
- Set different snare reverbs in different sections.

A mix can add or create interest.

With all of this in mind, it is worth remembering that not all music is meant to be attention-grabbing. Brian Eno once said about ambient music that it should be ignorable and interesting at the same time. Some genres, such as trance, are based on highly repetitive patterns that might call for more subtle changes.

Frequency domain

The frequency domain is probably the hardest aspect of mixing to master. Some approximate that half of the work that goes into a mix involves frequency treatment.

The frequency spectrum

Most people are familiar with the bass and treble controls found on many hi-fi systems. Most people also know that the bass control adjusts the low frequencies (lows) and the treble the high frequencies (highs). Some people are also aware that in between the lows and highs are the mid-frequencies (mids).

Our audible frequency range is 20 Hz to 20 kHz (20,000 Hz). It is often subdivided into four basic bands: lows, low-mids, high-mids, and highs. This division originates from the common four-band equalizers found on many analog desks. There is no standard as to the amount of bands or where exactly each band begins and ends, and on most equalizers the different bands overlap. Roughly speaking, the crossover points are at 250 Hz, 2 kHz, and 6 kHz, as illustrated on the frequency-response graph in Figure 7.1. Of the four bands, the extreme bands are easiest to recognize since they either open or seal the frequency spectrum. Having the ability to identify the lower or higher midrange can take a bit of practice.
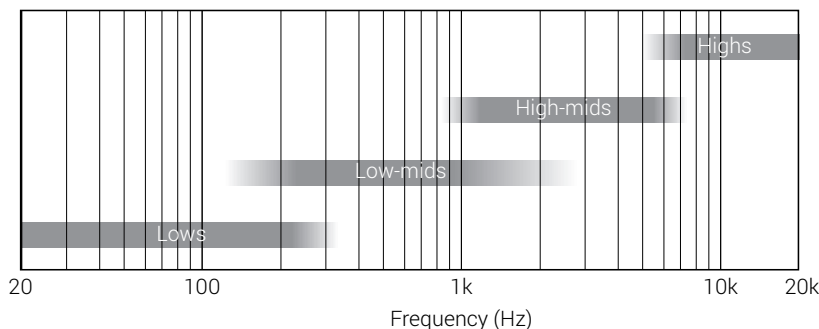
Figure 7.1 The four basic bands of the frequency spectrum.

Frequency balance and common problems

Achieving frequency balance (also referred to as tonal balance) is a prime challenge in most mixes. Here again, it is hard to define what is a good tonal balance, and our ears very quickly get used to different mix tonalities. In his book *Mastering Audio*, Bob Katz suggests that the tonal balance of a symphony orchestra can be used as a reference for many genres. Although there might not be an absolute reference, my experience shows that seasoned engineers have an unhesitating conception of frequency balance. Moreover, having a few engineers listening to an unfamiliar mix in the same room

would result in remarkably similar opinions. Although different tonal balances might all be approved, any of the issues presented below are rarely argued on a technical basis (although they might be argued on an artistic basis—"I agree it's a bit dull, but I think this is how this type of music should sound").

The following tracks demonstrate the isolation of each of the four basic frequency bands. All tracks were produced using an HPF, LPF, or a combination of both. The crossover frequencies between the bands were set to 250 Hz, 2 kHz, and 6 kHz. The filter slope was set to 24 dB/oct.

Track 7.1: Drums Source
The source, unprocessed track used in the following samples.

Track 7.2: Drums LF Only
Notice the trace of crashes in this track.

Track 7.3: Drums LMF Only

Track 7.4: Drums HMF Only

Track 7.5: Drums HF Only
And the same set of samples with vocals:

Track 7.6: Vocal Source

Track 7.7: Vocal LF Only

Track 7.8: Vocal LMF Only

Track 7.9: Vocal HMF Only

Track 7.10: Vocal HF Only

Plugin: Digidesign DigiRack EQ 3

The most common problems with frequency balance involve the extremes. A mix is boomy if there is an excess of low-frequency content, and thin if there is a deficiency. A mix is dull if there are not enough highs, and brittle if there are too many. Many novice engineers come up with mixes that present these problems, mainly due to their colored monitors and lack of experience in evaluating these ranges. Generally, the brighter the instruments are, the more defined they become and the more appealing they can sound. There is a dangerous tendency to brighten up everything and end up with very dominant highs. This is not always a product of equalization—enhancers and distortions also add high-frequency content. Either a good night's sleep or a comparison to a reference track will normally help pinpoint these problems. Overemphasized highs are also a problem during mastering, since mastering engineers can use some high-frequency headroom for enhancement purposes. A lightly dull mix can be made brighter with relative ease. A brittle mix

can be softened, but this might limit the mastering engineer when applying sonic enhancements.

Low-frequency issues are also prevalent due to the great variety of playback systems and their limited accuracy in reproducing low frequencies (such as in most bedroom studios). For this very reason, low frequencies are usually the hardest to stabilize, and it is always worth paying attention to this range when comparing a mix on different playback systems.

Track 7.11: Hero Balanced
Subject to taste, the monitors used, and the listening environment, this track presents a relatively balanced frequency spectrum. The following tracks are exaggerated examples of mixes with extreme excesses or deficiencies.

Track 7.12: Hero Lows Excess

Track 7.13: Hero Lows Deficiency

Track 7.14: Hero Highs Excess

Track 7.15: Hero Highs Deficiency

In general, it is more common for a mix to have an excess of uncontrolled low-end than to be deficient in it. The next most problematic frequency band is the low-mids, where most instruments have their fundamentals. Separation and definition in the mix is largely dependent on the work done in this busy area, which can very often be cluttered with nonessential content.

One question always worth asking is: Which instrument contributes to which part of the frequency spectrum? Muting the kick and the bass, which provide most of the low-frequency content, will cause most mixes to sound powerless and thin. Based on the *percussives weigh less* axiom, muting the bass is likely to cause more low-end deficiency. Yet, our ears might not be very discriminating when it comes to short absences of some frequency ranges—a hi-hat can appear to fill the high-frequency range, even when played slowly.

One key aspect of the frequency domain is **separation**. In our perception, we want each instrument to have a defined position and size on the frequency spectrum. It is always a good sign if we can say that, for instance, the bass is the lowest, followed by the kick, then the piano, snare, vocals, guitar, and cymbals. The order matters less—the important thing is that we can separate one instrument from the others. Then we can try to see whether there are empty areas: Is there a smooth frequency transition between the vocals and the hi-hats, or are they spaced apart with nothing between them? This type of evaluation, as demonstrated in Figure 7.2, is rather abstract. In practice, instruments can span the majority of the frequency spectrum. But we can still have a good sense of whether instruments overlap, and whether there are empty frequency areas.
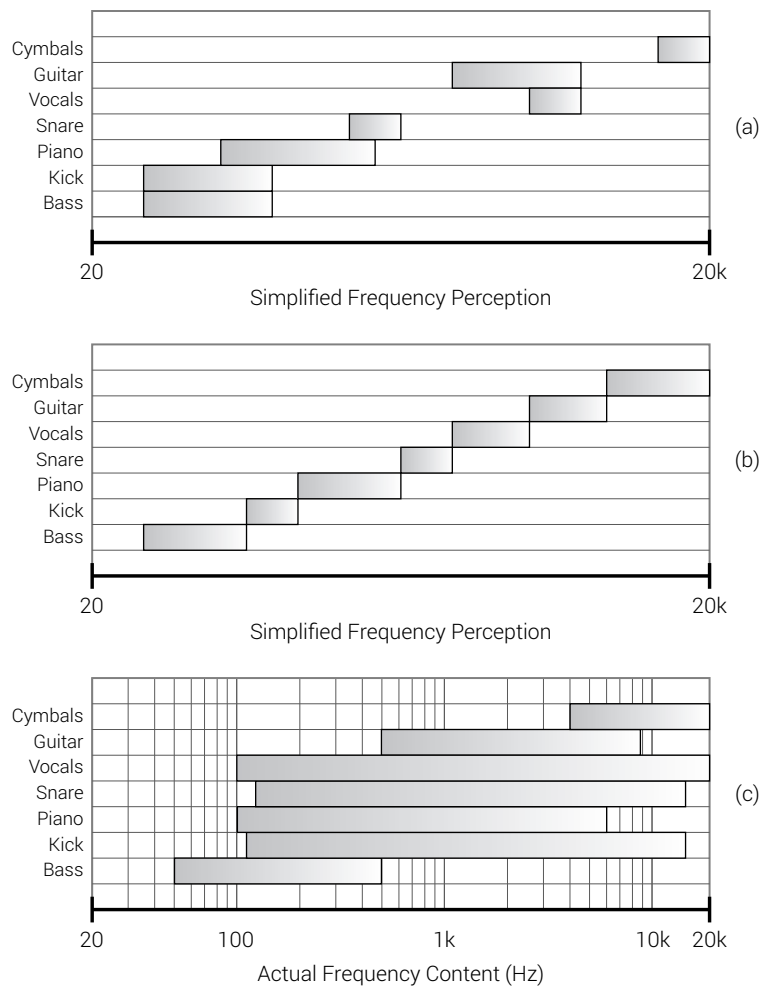
Figure 7.2 Instrument distribution on the frequency spectrum. (a) An unbalanced mix with some instruments covering others (no separation) and a few empty regions. (b) A balanced mix where each instrument has a well-defined range that does not overlap with other instruments. Altogether, the various instruments constitute a full, continuous frequency response. (c) The actual frequency ranges of the different instruments.

The frequency domain and other objectives

Since masking is a frequency affair, **definition** is bound to the way various instruments are crafted into the frequency spectrum. Fighting elements mask one another and have competing content on specific frequency ranges. For example, the bass might mask the kick since both have healthy low-frequency content; the kick's attack might be masked by the snare, as both appear in the high-mids. We equalize various instruments to increase their definition, a practice often done in relation to other masking instruments. Occasionally, we might also want to decrease the definition of instruments that stand out too much.

Frequency content relates to **mood**. Generally, low-frequency emphasis relates to a darker, more mysterious mood, while high-frequency content relates to happiness and

liveliness. Although power is usually linked to low frequencies, it can be achieved through all the areas of the frequency spectrum. There are many examples of how the equalization of individual instruments can help in conveying one mood or another; these are discussed later, in Chapter 15.

Many DJs are very good at balancing the frequency spectrum. An integral part of their job is making each track in the set similar in tonality to its preceding track. One very common DJ move involves sweeping up a resonant high-pass filter for a few bars and then switching it off. The momentary loss of low frequencies and the resonant high frequencies creates some tension and makes the reintroduction of the kick and bass very exciting for some clubbers, with those already excited sometimes turning overexcited. Mixing engineers do not go to the same extremes as DJs, but attenuating momentarily low frequencies during a drop or a transition section can achieve a similar effect. Some frequency **interest** occurs naturally with arrangement changes—choruses might be brighter than verses due to the addition of some instruments. Low frequencies usually remain intact despite these changes, but, whether mixing a recorded or an electronic track, it is possible to automate a shelving EQ to add more lows during the exciting sections.

Level domain

High-level signals can cause unwanted distortion—once we have cleared up this technical issue, our mix evaluation process is much more concerned with the **relative levels** between instruments rather than their absolute level (which varies in relation to the monitoring levels). Put another way, the question is not "How loud?" but "How loud compared with other instruments?" The latter question is the basis for most of our level decisions—we usually set the level of an instrument while comparing it to the level of another. There are huge margins for personal taste when it comes to relative levels (as your excerpt set from Chapter 1 should prove), and the only people who are likely to get it truly wrong are the fledglings—achieving a good relative level balance is a task that for most of us will come naturally.

However, as time goes by, we learn to appreciate how subtle adjustments can have a dramatic effect on the mix. For instance, even a boost of 2 dB on pads can make a mix much more appealing. In contrast to the widespread belief, setting relative levels involves much more than just moving faders; equalizers and compressors are employed to adjust the perceived loudness of instruments in more sophisticated ways—getting an exceptional relative level balance is an art and requires experience.

Levels and balance

It is worth discussing what a good relative level balance is. Some mixing engineers are experienced enough to create a mix where all the instruments sound as loud. But only a few, mostly sparse mixes might benefit from such an approach. All other mixes usually call for some variety. In fact, trying to make everything equally loud is often a self-defeating habit that novice engineers adopt—it can be both impractical and totally unsuitable. It is worth remembering that the relative level balance of a mastered mix is likely to be tighter—raw mixes usually have greater relative levels variety.

Track 7.16: Level Balance Original
This track presents a sensible relative level balance between the different instruments. The following tracks demonstrate variations of that balance, arguably for the worse.

Track 7.17: Level Balance 1 (Guitars Up)

Track 7.18: Level Balance 2 (Kick Snare Up)

Track 7.19: Level Balance 3 (Vocal Up)

Track 7.20: Level Balance 4 (Bass Up)

Setting relative balance between the various instruments is usually determined by their importance. For example, the kick in a dance track is more important than any pads. Vocals are usually the most important in rock and pop music. (A very common question is: "Is anything louder than the vocals?") Maintaining sensible relative levels usually involves gain-rides. As our song progresses, the importance of various instruments might change, as in the case of a lead guitar that is made louder during its solo section.

Steep level variations of the overall mix might also be a problem. These usually occur as a result of major arrangement changes that are very common in sections such as intros, breaks, and outros. If we do not automate levels, the overall mix level can dive or rise in a disturbing way. The level of the crunchy guitar on the first few seconds of Nirvana's "Smells Like Teen Spirit" was gain-ridden exactly for this purpose—having an overall balanced mix level. One commercial example of how disturbing the absence of such balance can be can be heard on "Everything for Free" by K's Choice. The mix just explodes after the intro and the level burst can easily make you jump off your seat. This track also provides, in my opinion, an example of vocals not being loud enough.

Levels and interest

Although notable level changes of the overall mix are undesirable, we want some degree of level variations in order to promote interest and reflect the intensity of the song. Even if we do not automate any levels, the arrangement of many productions will mean a quieter mix during the verse and a louder one during the chorus. Figure 7.3 shows the waveform of "Witness" by The Delgados, and we can clearly see the level variations between the different sections.

Very often, however, we spice up these inherent level changes with additional automation. While mastering engineers achieve this by automating the overall mix level, during mixdown we mostly automate individual instruments (or instrument groups). The options are endless: the kick, snare, or vocals might be brought up during the chorus; the overheads might be brought up on downbeats or whenever a crash hits (crashes often represent a quick intensity burst); we might bring down the level of reverbs to create a tighter, more focused ambiance; we can also
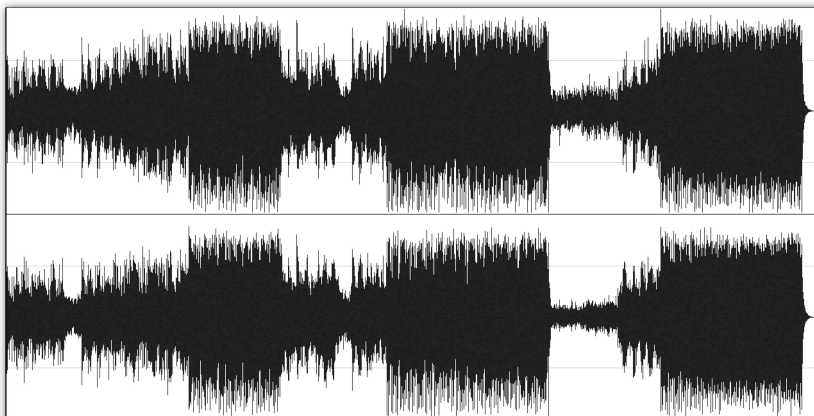
Figure 7.3 The level changes in this waveform clearly reveal the verses, choruses, and various breaks.

automate the drum mix in a sinusoidal fashion with respect to the rhythm—it has been done before.

Level automation is used to preserve overall level balance, but also to break it.

Levels, mood, and definition

If we consider a dance track, the level of the kick is tightly related to how well the music will affect us—a soft kick will fail to move people in a club. If a loud piano competes with the vocals in a jazz song, we might miss the emotional message of the lyrics or the beauty of the melody. We should always ask ourselves what the emotional function of each instrument is and how it can enhance or damage the mood of a song, then set the levels accordingly. In mixing, there are always alternatives—it is not always the first thing that comes to mind that works best. For example, the novice's way to create power and aggression in a mix is to have very loud, distorted guitars. But the same degree of power and aggression can be achieved using other strategies, which will enable us to bring down the level of the masking guitars and create a better mix altogether.

Definition and levels are linked: the louder an instrument is, the more defined it is in the mix. However, an instrument with frequency deficiencies might not benefit from a level boost—it might still be undefined, just louder. Remember that bringing up the level of a specific instrument can cause the loss in definition of another.

Dynamic processing

So far, we have discussed the overall level of the mix, and the relative levels of the individual instruments it consists of. Another very important aspect of a mix relates to noticeable level changes in the actual performance of each instrument (micromixing). Inexperienced musicians can often produce a single note, or a drum hit, that is either very loud or quiet compared with the rest of the performance. Rarely do vocalists

produce a level-even performance. These level variations, whether sudden or gradual, break the relative balance of a mix, and therefore we need to contain them using gain-riding, compression, or both. We can also encounter the opposite, mostly caused by over-compression, where the perceived level of a performance is too flat, making instruments sound lifeless and unnatural.

Track 7.21: Hero No Vocal Compression
In this track, the vocal compressor was bypassed. Not only is the level fluctuation of the vocal disturbing, but the relative level between the voice and other instruments alters.

Stereo domain

Stereo image criteria

The stereo panorama is the imaginary space between the left and right speakers. When we talk about stereo image, we either talk about the stereo image of the whole mix or the stereo image of individual instruments within it—a drum kit, for example. One incorrect assumption is that stereo image is only concerned with how far to the left or right instruments are panned. But stereo image involves slightly more advanced concepts, classified like so (also illustrated in Figure 7.4):

• **Localization**—concerned with where the sound appears to come from on the left–right axis.
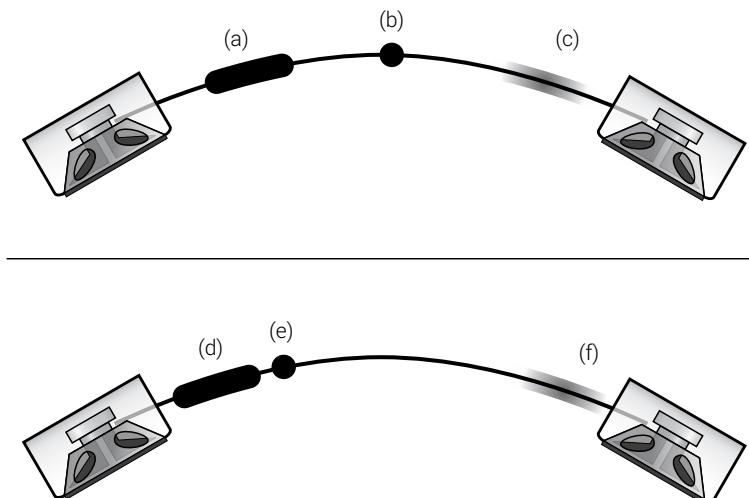
Figure 7.4 Stereo criteria. Localization is concerned with our ability to discern the exact position of an instrument. For example, the position of instrument (b)/(e). Stereo width is concerned with the length of stereo space the instrument occupies. For example, (a) might be a piano wider in image than a vocalist (b). If we can't localize precisely the position of an instrument, its image is said to be smeared (or diffused), like that of (c). The difference between (a, b, c) and (d, e, f) is the way they are spread across the stereo panorama.

- **Stereo width**—how much of the stereo image the sound occupies. A drum kit can appear narrow or wide, as can a snare reverb.
- **Stereo focus**—how focused the sounds are. A snare can appear to be emanating from a very distinct point in the stereo image, or can be unfocused, smeared, and appear to be coming from "somewhere over there."
- **Stereo spread**—how exactly the various elements are spread across the stereo image. For example, the individual drums on an overhead recording can appear to be coming mostly from the left and right, and less from the center.

To control these aspects of the stereo panorama, we use various processors and stereo effects, mainly pan-pots and reverbs. As we will learn in Part II, stereo effects such as reverbs can fill gaps in the stereo panorama, but can also damage focus and localization.

Stereo balance

Balance is the main objective in the stereo domain. Above all, we are interested in the **balance between the left and right** areas of the mix. Having a stereo image that shifts to one side is unpleasant and can even cause problems if the mix is later cut to vinyl. Mostly, we consider the **level balance** between left and right—if one channel has more instruments panned to it, and if these instruments are louder than the ones panned to the other channel, the mix image will converge on one speaker. It is like having a play with most of the action happening stage left. Level imbalance between the left and right speakers is not very common in mixes, since we have a natural tendency to balance this aspect of the mix. It is worth remembering that image shifting can also happen at specific points throughout the song, for example when a new instrument is introduced. One tool that can help us in identifying image shifting is the L/R swap switch.

There are more than a few commercial mixes that involve stereo imbalance. One example is "Give It Away" by Red Hot Chili Peppers, where John Frusciante's guitar appears nearly fully to the right, with nothing to balance the left side of the mix. The cymbals panned to the left on Kruder & Dorfmeister's remix for "Useless" are another example.

A slightly more common type of problem is **stereo frequency imbalance**. While the frequency balance between the two speakers is rarely identical, too much variation can result in image shifting. Even two viola parts played an octave apart can cause image shifting if panned symmetrically left and right (the higher-octave part will draw more attention). Stereo frequency balance is much to do with arrangement—if there is only one track of an instrument with a very distinguished frequency content, it can cause imbalance when panned. We either pan such an instrument more toward the center (to minimize the resultant image shifting) or we employ a stereo effect to fill the other side of the panorama. Hi-hats are known for causing such problems, especially in recorded productions, where it makes less sense to add delay to them.

While left/right imbalance is the most noticeable of all stereo problems, **stereo spread imbalance** is the most common. In most mixes, we expect to have the elements spread across the stereo panorama so there are no lacking areas, like a choir is organized onstage. The most obvious type of stereo spread imbalance is a nearly monophonic mix—one that makes very little use of the sides, also known as an *I-mix* (Figure 7.5). While such a mix can be disturbing, it can sometimes be appropriate. Nearly monophonic mixes are most prevalent in hip-hop—the beat, bass, and vocals are gathered around the center, and only a few sounds, such as reverbs or backing pads, are sent to the sides. One example of such a mix is "Apocalypse" by Wyclef Jean, although not all hip-hop tracks are mixed that way.

Figure 7.5 An I-mix. The dark area between the two speakers indicates the intensity area. The white area indicates the empty area. An I-mix has most of its intensity around the center of the stereo panorama, with nothing or very little toward the extremes.

Another possible issue with stereo spread involves a mix that has a weak center, and most of its intensity is panned to the extremes (it is therefore called a *V-mix*) (Figure 7.6). V-mixes are usually the outcome of a creative stereo stratagem (or the result of three-state pan switches such as those found on very old consoles); otherwise, they are very rare.
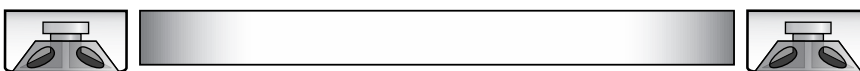
Figure 7.6 A V-mix. The term describes a mix that has very little around the center, and most of the intensity is spread to the extremes.

The next type of stereo spread imbalance, and a very common one, is a combination of the previous two, known as the *W-mix*—a mix that has most of its elements panned hard-left, center, and hard-right (Figure 7.7). Many novices produce such mixes because they tend to pan every stereo signal to the extremes. A W-mix is not only unpleasant, but in a dense arrangement can be considered a waste of stereo space. Here again, the arrangement plays a role—if there are few instruments, we need to widen them to fill spaces in the mix, but this is not always appropriate. An example of a W-mix is the verse sections of "Hey Ya!" by OutKast; during the chorus, the stereo panorama becomes more balanced, and the change between the two is an example of how interest can originate from stereo balance variations.
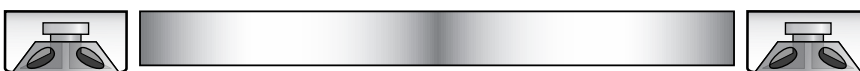
Figure 7.7 A W-mix involves intensity around the extremes and center, but not much in between.

The last type of stereo spread imbalance involves a stereo panorama that is lacking in a specific area, say between 13:00 and 14:00. It is like having three adjacent players missing from a row of trumpet players. Figure 7.8 illustrates this. Identifying such a problem is easier with a wide and accurate stereo setup. This type of problem is usually solved with panning adjustments.

Figure 7.8 A mix lacking a specific area on the stereo panorama. The empty area is seen next to the right speaker.

I- and W-mixes are the most common types of stereo spread imbalance.

Track 7.22: Full Stereo Spread
The full span of the stereo panorama has been utilized in this track. The following tracks demonstrate the stereo spread issues discussed above.

Track 7.23: W-Mix

Track 7.24: I-Mix

Track 7.25: V-Mix

Track 7.26: Right Hole
In this track, there is an empty area on the right-hand side of the mix.

Stereo image and other objectives

Although it might seem unlikely, stereo image can also promote various **moods**. Soundscape music, chill-out, ambient, and the like tend to have a wider and less-focused stereo image compared with less relaxed genres. The same old rule applies—the less natural, the more powerful. A drum kit as wide as the stereo panorama may sound more natural than a drum kit panned around the center. In "Witness" by The Delgados, Dave Fridmann chose to have a wide drum image for the verse, and then a narrow image during the more powerful chorus.

We have seen a few examples of how the stereo image can be manipulated in order to achieve **interest**. Beethoven achieved this by shifting identical phrases across the orchestra, while in "Hey Ya!" the stereo spread changes between the verse and the chorus. It is worth listening to Nirvana's "Smells Like Teen Spirit" and noting the panning strategy and the change in stereo image between the interlude and the chorus. Other techniques include panning automation, auto-pans, stereo delays, and so forth. We will see in Chapter 14 how panning can also improve **definition**.

Depth

For a few people, the fact that a mix involves a front/back perspective comes as a surprise. Various mixing tools, notably reverbs, enable us to create a sense of depth in our mix—a vital extension to our sound stage, and our ability to position instruments within it.

All depth considerations are relative. We never talk in terms of how many meters away—we talk in terms of in front of or behind another instrument. The depth axis of our mix starts with the closest sound and finishes with the farthest away. A mix where instruments are very close can be considered tight, and a mix that has an extended depth can be considered spacious. This is used to reflect the **mood** of the musical piece.

The depth field is an outsider when it comes to our standard mixing objectives. We rarely talk about a balanced depth field since we do not aim to have our instruments equally spaced in this domain. A **coherent depth** field is a much more likely objective. Also, in most cases, the farther away an instrument is, the less defined it becomes; crafting the depth field while retaining the definition of the individual instruments can be a challenge. A classical concert with the musicians walking back and forth around the stage would be chaotic. Likewise, depth variations are uncommon, and usually instruments move back and forth in the mix as a creative effect, or in order to promote importance (just like a trumpet player would walk to the front of the stage during his or her solo).

Track 7.27: Depth Demo
This track involves three instruments in a spatial arrangement. The lead is dry and foremost, the congas are positioned behind it, and the flute-like synth is placed way at the back. The depth perception in this track should be maintained whether listening on the central plan between the speakers, at random points in the room, or even outside the door.

Track 7.28: No Depth Demo
This track is the same as the previous arrangement, but excludes the mix elements that contributed to depth perception.

Plugin: Audio Ease Altiverb
Percussion: Toontrack EZdrummer

We are used to sonic depth in nature and we want our mixes to re-create this sense. The main objective is to create something natural or otherwise artificial but appealing. Our decisions with regard to instrument placement are largely determined by importance and what we are familiar with in nature. We expect the vocals in most rock productions to be the closest, just like a singer is frontmost on stage. But in many electronic dance tracks, the kick will be in front and the vocals will appear behind it,

sometimes even way behind it. Then we might want all drum kit components to be relatively close on the depth field because this is how they are organized in real life.

Every two-speaker system has an integral depth, and sounds coming from the sides appear closer than sounds coming from the center. As demonstrated in Figure 7.9, the bending of the front image changes between speaker models, and when we toggle between two sets of speakers our whole mix might appear to shift forward or backward. The sound stage created by a stereo setup has the shape of a skewed rectangle and its limits are set to the angle between our head and the speakers (although we can create an out-of-speaker effect using phase tricks). As illustrated in Figure 7.10, the wider the angle between our head and the speakers, the wider the sound stage will become, especially for distant instruments.
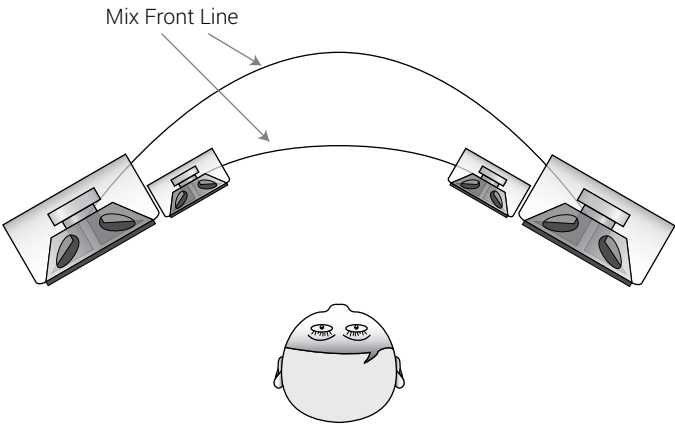
Figure 7.9 The perceived frontline of the mix appears in the stereo image as an arch, the shape of which varies between speaker models.
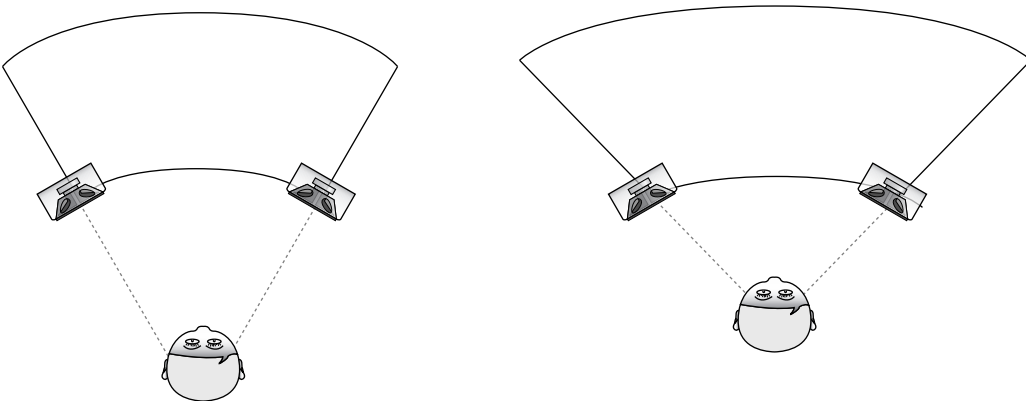
Figure 7.10 The width of the sound stage will change as our head moves forward or backward.

These two tracks demonstrate the integral depth of a two-speaker system. Pink noise and a synth are swept in a cyclic fashion between left and right. When listening on the central plane between properly positioned monitors, sounds should appear to be closest when they reach the extremes, and farther away as they approach the center. It should also be possible to imagine the actual depth curve. Listening to these tracks with eyes shut might make the apparent locations easier to discern.

Track 7.29: Swept Pink Noise

Track 7.30: Swept Synth

These tracks were produced using Logic, with –3 dB pan law (which will be explained later).

Part II

Tools

8 Monitoring

Studio loudspeakers are called *monitors*. This chapter opens Part II for good reason—a good monitoring environment is an absolute mixing requisite, and should be high up in budget planning, whether you mix at home or in a professional studio. The main conclusion of this chapter is worth revealing already—the monitors are just one variable; it is the monitors, their position, the listener's position, and the acoustic properties of the room that dictate the overall quality of the monitoring environment. Having great monitors badly positioned in a problematic room is like having a Ferrari that only goes up to second gear.

How did we get here?

Sound reproduction

In order to reproduce sound, loudspeaker drivers displace air in response to an incoming voltage that corresponds to a waveform. There is a fundamental difference in the way low and high frequencies are reproduced. Low frequencies call for a rigid, big cone that is capable of displacing a large mass of air. Conversely, high frequencies require a light and small diaphragm that can move rapidly. The two requirements obviously conflict. Additionally, low frequencies require a large amount of excursion compared with high frequencies. If a single driver produces both low and high frequencies, the cone displacement caused by low frequencies results in unwanted phase shifts for the high frequencies.

For this reason, the reproduction of the full audible range by a single driver yields quality below professional standards. Therefore, loudspeakers utilize two or more drivers, which are referred to as two-way design, three-way design, and so forth. We must make sure that each driver is only fed with the frequencies it specializes in reproducing. A device called a *crossover* is used to split the incoming signal into different frequency bands (Figure 8.1). A typical two-way loudspeaker would have its crossover around 2 kHz, sending each band to a different driver. It is impossible to build a crossover that simply slices the frequency spectrum in a brick-wall fashion, so there is always some overlapping bandwidth where identical frequencies are sent to both drivers. The insertion of a crossover into the signal path introduces various issues that color the sound; manufacturers address them in different ways, but no system is perfect.

One issue with multiway design is that the complete frequency spectrum is produced from different points in space. This might cause unwanted phase interaction when the
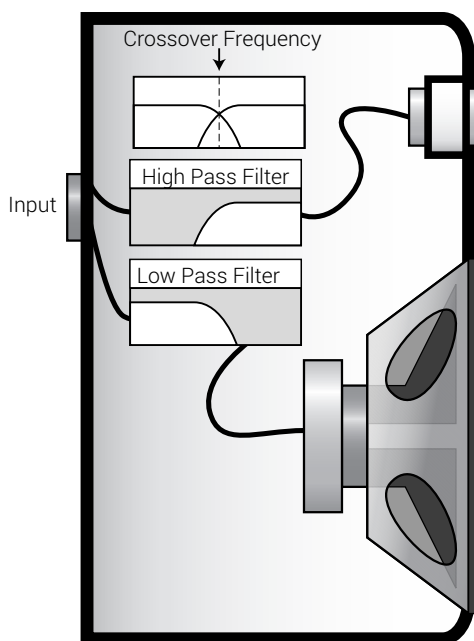
Figure 8.1 Crossover network in a two-way loudspeaker. As the signal enters the speaker, a filter network splits it into two different bands. The low frequencies are sent to the woofer, the high ones to the tweeter.

sounds emitted from the different drivers are summed acoustically. To solve this, some manufacturers (such as Tannoy) have designs where the high-frequency driver is fitted into the center of the low-frequency driver (where the dust dome is normally present).

Two-way studio monitors are able to reproduce the audible frequency spectrum with an acceptable quality. A design that involves more than two drivers will bring about better quality only if the problems introduced by the additional drivers and crossovers are addressed. Naturally, a design that addresses these problems involves a higher price tag. A two-way studio monitor, therefore, is more likely to be a better buy than a domestic three-way loudspeaker of the same price.

All ordinary studio monitors can produce frequencies up to 20 kHz. The lower limit of the frequency response is largely determined by the size of the woofer: 6″ and 8″ are very common diameters. (The size often forms part of the model name, although it rarely denotes the exact diameter.) These set the lower frequency limit to around 55 and 45 Hz, respectively. A loudspeaker still produces frequencies below these quoted limits, although these are gradually rolled off.

Auratones, near-fields, and full-range monitors

In the early days of mixing, the assumption was that engineers needed the best monitoring system available. But the main monitors (often referred to as "mains") in commercial studios were far superior to those consumers actually used, so many mixes did not translate well in consumer environments. Mixing engineers soon realized that they needed some speakers that imitated the sound of real-world consumer loudspeakers,

Figure 8.2 The Auratones 5C are the mini-speakers mounted on the meter bridge. Behind them are the Genelec 1031s near-fields.
Source: Courtesy of SAE, London.

and the Auratones 5C, also known as "the cubes," did the trick (Figure 8.2). These small, and not surprisingly cubic, single-driver speakers had a defined midrange, and they got the nickname "horrortones" because of their unappealing sound. Nevertheless, engineers soon realized that these small speakers could be used for more than simulating the sound of cheap systems.

Being found in many studios even today, the Auratones are very often used for critical level adjustments (notably vocals), and to evaluate definition. (For example, we can tell there's a problem if the kick and bass are lost with these monitors selected.) Having pronounced mids, they also help us to tidy up and equalize this area of the mix, where most instruments have their fundamentals and lower harmonics.

But something was still missing between the small Auratones and the big mains, and so a new type of compact monitor was released to fill the gap—the near-field monitors. The acoustical term "near-field" is misleading; a more apt term would be "close-field" monitors, which better describes their position—within the critical distance. Very often, these monitors are placed on top of the console's meter bridge or on stands right behind it. The vast majority of mixes are done using near-fields. Even in professional studios, where mains exist, mixing engineers might use near-fields for the majority of the mixing process, and only occasionally refer to a different set of monitors. The most common type of near-field nowadays is the active, 8″ two-way monitor.

The NS10s

No book about mixing would be complete without discussing the Yamaha NS10s. In 1978, Yamaha released a bookshelf loudspeaker intended for home use called the NS10M. It only took a few years before most music studios had a pair installed. Many people tried to explain the immense popularity of the NS10s. Some said that they were the natural successors to the Auratones, providing a compact speaker, with better sound but not too flattering. Many engineers testified that the midrange definition of these speakers and their tight low-end was highly beneficial for rock and pop mixes. Others declared that if something sounded right on the NS10s, it would translate well on most consumer systems. Another hypothesis is that their distinctive white cones had big shares in propelling them to a cult status. In his book, *Recording Studio Design*,

Philip Newell dedicated several pages to the NS10s and presented detailed research aiming to solve the mystery (including a crane-mounted SSL on which the speakers were then fixed). While some engineers wouldn't touch them with a 10-foot pole, they have become standard in the audio-engineering field, particularly in mixing.

The studio popularity of the NS10s caught Yamaha by surprise. They never designed these speakers for professional use and it soon became apparent that the classic version (NS10M) had some issues when used in studio environments. First, their vertical mounting meant they often obscured the main monitors. Second, they were unable to withstand the abusive studio levels and the tweeters often blew. Finally, and most famously, they had harsh and emphasized highs, which led many people to cover the tweeters with tissue paper. One enthusiastic engineer, Bob Hodas, even researched what type of tissue to use, how many layers, and how far they should be from the tweeters. Yamaha themselves used tissue paper as part of their research while developing a re-designed model, and in 1987 the Yamaha NS10M Studio were released. These were intended for horizontal mounting (and the direction of the label changed accordingly), had softened highs, and more robust tweeters; the cones remained white. As the material used for producing those white cones became unavailable, Yamaha discontinued the production of the NS10s in 2001. Like the Auratones, the NS10s are still found in many studios today, and there is still demand for them in the secondhand market.

Going full-range

While the limited bass response of near-fields might not be an issue for some mixes, it is crucial to have an extended low-frequency response in genres such as hip-hop, reggae, dance, and others. **Full-range** monitors are so-called because they can reproduce the full audible range from 20 Hz to 20 kHz. They are usually large, high-fidelity monitors that provide higher resolution compared with near-fields. In many studios, these monitors are flush mounted, which enhances their acoustic interaction with the room. Where these exist, we refer to them in order to check and stabilize the low-end of the mix. When switching to the full-range monitors, we hope for an

Figure 8.3 The Angel Two Studio at Abbey Road Institute, London. The three types of monitors can be seen here: the full-range are the Bower and Wilkins 802, the near-fields are the PMC twotwo.6 and the mini-speaker is the Avantone MixCube.

extended low-end without losing level balance. They are also useful when creating separation between the bass and the kick and, due to their high resolution, we also refer to them for refinements, such as subtle vocal equalization. The high quality of large full-range monitors makes them the favorite choice for mastering, classical music production, and many recording situations. But their sound is so far superior to anything that most listeners have access to that mixing engineers generally favor the smaller, less impressive, yet accurate near-fields.

Subwoofers

Space and budget limitations make full-range monitors a rare breed in home studios, but a dedicated subwoofer provides an alternative. Many professional monitor manufacturers offer a matching subwoofer to their near-fields range, which normally covers the 20–150 Hz range. An optimum configuration for such a setup often involves feeding the stereo mix into the subwoofer first. Most subwoofers have a built-in crossover that splits the frequency spectrum into two bands—the very low frequencies are summed to mono and sent to the subwoofer driver, and all other frequencies are sent to the near-fields through a dedicated stereo output at the subwoofer's rear.

Choosing monitors

Active vs. passive

The low-level, low-power line output of our desk or computer needs to be amplified to the powerful speaker-level in order for the mechanical components in a loudspeaker to move. There is always a driving amplifier in the signal chain prior to the loudspeaker drivers, and this amplifier can be either built into the cabinet or live as an external unit. A loudspeaker with no integrated amplifier is known as a **passive speaker**, and must be fed with a speaker-level signal that has been amplified by an external amplifier. Most multiway speakers contain a crossover within their cabinet. In the case of passive speakers, the crossover is a passive one (there are no active components) and is designed to operate at speaker level. External amplifiers have a huge influence on the overall sound. The NS10s, for example, can sound distinctively different when driven by different makes with different power ratings. There is some critical interaction between an amplifier and a loudspeaker—it is, essentially, one system—and the amplifier determines many aspects of the overall sound, such as transient response, low-frequencies reproduction, and distortion. Matching an amplifier to a loudspeaker is never an easy affair.

When connecting an amplifier to passive speakers, the following should be adhered to. Cables must be as short as possible and of equal length (for both the left and right speakers); if they are not of equal length, there will be stereo imbalance between the two speakers. It is important to make sure that the terminals are correctly connected with plus to plus and minus to minus; if the terminals are crossed on both speakers, the two channels will be in phase, but the overall mix will be phase-inverted. This means that the speaker cones will move backwards instead of forward, and vice versa, and might affect the sound. (This is based on the idea that cone suspension does not behave identically between front and back movements. As part of quality control, some mastering engineers check their masters with both speakers inverted.) An easy way to diagnose such incorrect cabling is when a kick pulls the cone rather than pushes it. If only one speaker is

cross-connected, then the left and right speakers are phase-inverted, which creates an extremely unfocused, out-of-speakers sonic image that makes accurate mixing an impossible task. Conventionally, the positive (plus) terminals are connected using the red wire (as in the audio proverb: "red is ring, right, and positive").

Loudspeakers with built-in amplifiers are called **powered speakers**. These have line-level inputs, commonly balanced XLR or ¼". They involve two common designs: the first involves a single amplifier followed by a passive crossover; the second involves an active crossover followed by two or more amplifiers—one for each driver. A speaker that has an active crossover is called an **active speaker**, and if there is more then one built-in amplifier it is called a **bi-amped** or **tri-amped**, depending on the number of amplifiers. The majority of powered studio monitors are active and multi-amped.

Active monitors are often shielded, which drains magnetic interference that might potentially occur between the speaker and other electronic devices such as old CRT screens. Some newer models also have digital inputs. While placing an A/D converter within a loudspeaker might seem to make as much sense as placing a blender in a microwave, such a practice prevents analog interference and minimizes the chance of ground loops.

There are not many cases in mixing where one thing has such a clear advantage over something else, but active speakers provide many advantages over passive ones. The fact that a speaker has a built-in amplifier removes the guesswork of matching an amplifier to a loudspeaker and leaves this to the professional designers. Manufacturers can fine-tune the performance of each component for optimal results, and usually the outcome is more cost-effective. Many active speakers include protection circuitry (built-in limiters) that makes them resistant to abuse. If we ignore room acoustics, identical models of active speakers installed in different places would sound highly consistent since they are always driven by the same built-in amplifiers. But, for all its advantages, there is no guarantee that an active monitor will perform better than a passive one. Although active speakers continue to gain popularity, passive monitors are still manufactured, some receiving much acclaim. With all of its disadvantages, perhaps the one benefit of a passive setup is increased configurability—the ability to mix and match amplifiers and speakers for truly fine-tuned results—something the serious techie takes pleasure in.

Enclosure designs and specifications

You may already be aware that some studio monitors have holes in their enclosure (called "ports," "vents," or "ducts" in professional jargon). Such a design is known as **dipole design** and includes sub-designs such as *vented enclosure* (e.g., Genelec 1031 A), *bass reflex* (e.g., Dynaudio BM 6 A), and *transmission line* (e.g., PMC TB2S-A). Designs with no ports have the air within the enclosure sealed, and are known as **monopole designs**. The most common monopole designs are either *sealed enclosure* (e.g., Yamaha NS10s) or *auxiliary bass radiator*, which is also known as *ABR* or *passive radiator* (e.g., Mackie HR824). While the concept behind the different designs is a fascinating one, it teaches us nothing about the fidelity of the final product—none of the designs ensure a better quality than the others.

There are, however, a few things worth mentioning. Dipole designs are more efficient and provide an extended low-frequency response compared with monopole designs of

the same size (although the low-end extension is potentially unfaithful and an inaccurate one). Conversely, monopole designs provide better damping of the woofer cone when the input signal dies abruptly. After a gated kick dies out, momentum will keep the woofer cone moving for a while before it comes to a halt; this extraneous movement generates unwanted low frequencies. While the woofer in monopole designs comes to a halt very quickly, waterfall plots show that some dipole monitors can produce frequencies lower than 100 Hz for more than 100 ms after the input dies. Sound-wise, monopole designs are said to deliver a tighter bass response, and it is no wonder that many professional subwoofers employ such a design.

Like many audio devices, the specification sheets for monitors are often inconsistent and can be misleading, teaching us very little about the actual quality of the product. Technical measurements such as signal-to-noise ratio, harmonic distortion, and max-imum output level are often dependent on the system used during measurements, and can be easily manipulated in favor of the manufacturer. The monitors' frequency-response graph, which depicts the dips and bumps of various frequencies across the spectrum, has no bearing on the perceived quality or accuracy of the speaker. A speaker with a relatively flat frequency response does not guarantee a better quality than a speaker with noticeable wiggles. Of course, one specification that we do care about is the quoted frequency range of the speaker, specifically its lower limit. There is an important difference between a speaker that rolls off low frequencies at 70 Hz and one that does so at 50 Hz. Here again, manufacturers can manipulate the specifications to show better figures.

Frequency response specifications are often given with a quoted error range, for example 58 Hz–20 kHz (±2 dB). A frequency response within a ±2 dB range means wiggles are no more than 2 dB from the average. The smaller the error range, the better, as ±10 dB means that there is normally at least one wiggle that is 10 dB away from the average response. Wiggles aside, higher error range is sometimes used to "extend" the lower frequency limit of a speaker on specification sheets. This is due to the fact that the lower-frequency limit is often determined by the lowest frequency in the lows roll-off range that is outside that error range. So the same speaker will have an extended lower-frequency specification with an error range of ±10 dB compared to that of ±2 dB. An error range of ±2 dB typically means the manufacturer does not take advantage of this range to extend the frequency range; ±10 dB should be treated as suspicious; when no error range is provided, the frequency response specification is generally useless.

A choice of experience

It should be clear by this point that monitors should not be selected based on their design or specifications, but on their actual sound. To be sure, one thing we do *not* want from our monitors is a flattering sound—we want them to reveal problems rather than conceal them. Accuracy and detail are the key qualities we are after. "Because they sound good" is the poorest reason to favor one brand over another. Unfortunately, many buyers fail to comprehend this. Even worse, some manufacturers will sacrifice quality in order to make their monitors more appealing in quick-listening tests. Some retailers have a showroom containing a clutter of many brands and models in an arrangement that bears no resemblance to the actual positioning of monitors in a real studio setup. This supermarket

approach does the monitors no justice whatsoever, as often their sound can be greatly influenced by their placement. It also often means that we are unable to assess other critical aspects, such as stereo image.

By way of analogy, buying monitors is like buying a mattress—no matter how good it looks in the shop or how sophisticated the specifications are, we can only tell how comfortable a mattress is after using it for a few days. Mixing engineers will usually become accustomed to their monitors after a while—and being used to them, will remain loyal and depend on them. There truly isn't a "magic model"—what one praises the other dislikes. Even the NS10s have been the subject of debate. As a guideline: the higher the price tag, *usually*, the higher the quality one should expect.

The room factor

No pair of speakers sounds the same, unless placed in the same room.

It may seem unreasonable to expend as much money on acoustic treatment as on monitors. But the truth is that an expensive set of high-fidelity monitors is pointless if deficiencies in room acoustics are not treated. We have mentioned the frequency-response graphs and waterfall plots of a loudspeaker. These measurements are taken by manufacturers in anechoic chambers, where the room itself is not a variable. This is only fair, as it would be unreasonable for each manufacturer to use a different room—each with its own unique effect. If manufacturers were to record their measurements in domestic rooms, their results would yield variations that could be six times worse than the anechoic measurements. In practice, these untaken worse measurements are the ones our rooms produce—to this extent or another.

Professional mixing facilities are designed by specialists; therefore, this section will only cover the most relevant aspects for smaller project studios. A complete discussion of all the acoustic factors that affect our monitoring environment is far beyond the scope of this book, but the most important factors for our purposes are briefly covered below.

Room modes

Room modes are discussed in great detail in many books. A full exploration of room modes is long, technical, and requires some background knowledge. Below is a short and simplified explanation of this topic.

Low frequencies propagate in a spherical fashion. For the sake of simplicity, imagine that low frequencies emitted from our monitors travel equally in all directions. Also, whenever low frequencies come into contact with a surface, imagine that a new sound source is created at the point of contact, as if a new speaker has been placed there.

The simplest way of describing standing waves involves the use of two parallel walls. Sound emitted from a speaker will hit the left wall, bounce back to the right wall, and then bounce back and forth between the two. We can think of the sound as trapped between the two walls. Since every time sound hits a surface some of its energy is

being absorbed, after a while the sound dies out. However, continuous monitor output might constantly reinforce waves already trapped.

If the frequency of the trapped waves is mathematically related to the length of the trap, a predictable interaction between all the trapped waves will cause that frequency to be either attenuated or boosted at different points along the trap. For example, halfway between the two walls that frequency might be barely audible, and next to the wall that frequency might be overemphasized. Waves with these characteristics are called **standing waves** and a problematic frequency is described as a resonant **room mode**. The laws of acoustics say that if a specific frequency is trapped in a room, all of its harmonics are also trapped. For example, if 50 Hz is the lowest resonant mode, then 100, 150, 200, 250 Hz, and so forth will also be trapped.

If traps could only form between two parallel surfaces, each room would only have a relatively small set of three problematic frequencies (one for each dimension) and their harmonics. However, traps can also form between four or six surfaces, which results in a complex set of room modes. If two room dimensions are identical, then the problem is twofold since the same frequency is trapped along two dimensions. Cubic rooms are the worst since all three dimensions are identical. Rooms with dimensions that share common multiples (e.g., 3 × 6 × 9 m) are also very problematic. Despite the common belief, non-parallel walls do not rectify the standing waves problem—they only make the room modes' distribution more complex.

The lowest resonant frequencies are also the most profound ones. In small rooms, these lower frequencies are well within our audible range. The formula to calculate these fundamental frequencies is quite simple: $f = 172 / d$ (d refers to the dimension in meters). So, for example, a dimension of approximately 3 m will cause a fundamental room mode at 57 Hz. The bigger the room, the lower the fundamental resonant frequency. For example, a dimension of 8 m will cause a room mode at 21.5 Hz, which is less critical than 57 Hz for mixing applications. Also, as we climb up the frequency scale, the increased density of room modes makes them less profound, and around 500 Hz we can disregard them. In bigger rooms, problems start lower on the frequency scale, and also end lower—room modes might no longer be an issue above 300 Hz. Big rooms are therefore favored for critical listening applications such as mixing or mastering.

It is crucial to understand that room modes always cause problems at the same frequencies, but the problems are not consistent throughout the room. Eventually, each point in the room has its own frequency response. Room mode calculators are freely available over the Internet—they will output a graph showing problematic frequencies given the room dimensions. However, they do not tell us where exactly in the room we should expect these problems to be noticeable, or in other words what the frequency response of each point in the room will be.

Luckily, we can quite easily perform a practical test that teaches us just that. It involves playing a sine wave and comparing different frequencies between 20 and 600 Hz. Obviously, we usually listen at our mixing position, but moving around the room would demonstrate how drastically one frequency can be attenuated or boosted at different points across the room. The results of this experiment can be quite surprising for people trying it for the first time—you might learn that a specific

frequency (say, 72 Hz) is inaudible when you sit but clearly heard when you stand; you might also learn that one frequency (say, 178 Hz) is noticeably louder than a nearby frequency (say, 172 Hz).

One more aspect of room modes worth discussing is that they affect the speakers' ability to reproduce problematic frequencies. If a speaker is positioned at a point where standing waves cause a specific frequency boost, the speaker will be able to produce more of that frequency. The opposite case only applies to monopole speakers—if the speaker is located at a point where a specific frequency is attenuated, the driver will have difficulties producing that frequency.

Provided on the website are a few 30-second-long test tones for readers to experiment with. While listening to each frequency, move around your room to see how at different points that specific frequency is boosted or attenuated. Most readers trying this in a domestic room should recognize at least one frequency that suffers noticeable level variations across the room, which is the consequence of room modes.

Track 8.1: Sine 60 Hz

Track 8.2: Sine 65 Hz

Track 8.3: Sine 70 Hz

Track 8.4: Sine 75 Hz

Track 8.5: Sine 80 Hz

Track 8.6: Sine 90 Hz

Track 8.7: Pink Noise
Pink noise provides equal energy per octave and therefore is commonly used for acoustic measurement. While listening to this noise when seated in the listening position used for mixing, readers are encouraged to move their head back and forth, then sideways. Room modes might cause variation in low frequencies, while early reflections and the directivity of the tweeters might cause variations in high frequencies. One characteristic of a well-tuned room is that moving your head while listening to the music (or pink noise, for that matter) hardly alters the perceived frequency content.

Treating room modes

You might instinctively think that you could compensate for room mode effects by inserting a graphic equalizer before the monitors. If 172 Hz is attenuated by 6 dB, why not boost that frequency by 6 dB? When it comes to rectifying the acoustic response of the room, monitor equalization is considered futile and in most cases will do more harm than good. There are a few reasons for this. First is the simple fact that any equalization process has its own quality penalty, especially when using less than high-end equalizers (a high-precision graphical EQ can easily exceed the cost of acoustic treatment, which will yield better results). Second, the room response varies in relation to different positions within the room. It would make sense to compensate for audible problems at

the listening position, but such a treatment can cause greater problems at other positions in the room, including the points where the monitors are situated. Last, there is a difference between the way long sounds and transients excite room modes. Compared to the sound of a kick, room modes would greatly affect the sustained note of a bass guitar. Equalizing the monitors to make the bass sound good might make the kick sound boomy.

It is worth remembering what exactly we are trying to fix. The overall response at different places in a room is dependent on both the direct sound and the reflected sound. It is room modes caused by reflections that make a perfectly balanced mix unbalanced. But while trying to rectify room modes, monitor equalization also affects the direct sound, which can represent a well-balanced mix. There are situations where monitor equalization is appropriate, but these only happen when the direct sound itself experiences frequency alterations. For example, placing a loudspeaker next to the wall causes low-end emphasis.

To compensate for this, many active monitors offer switches for different degrees of bass roll-off. But this type of equalization is not intended to correct room modes; it is merely concerned with correcting coloration of the direct sound.

It should be clear by now that in order to treat room modes, we need to treat the reflections. This is achieved by two acoustic concepts: diffusion and absorption. Diffusers scatter sound energy, including the low-frequency energy of standing waves. Absorbers soak up sound energy. Diffusers are less welcome in small rooms, partly due to the fact that, in close proximity to the listening position, they can sometimes impair the overall response rather than enhance it. Absorbers, on the other hand, are a very practical solution. The idea is simple: if we absorb sound as it hits the wall, we damp the reflected energy, and therefore minimize the effect of standing waves. In anechoic chambers, there are no reflections and therefore no standing waves, but the unnatural response of these reflection-free spaces makes them highly unsuitable for mixing, and some people even find them unbearable. In most mixing situations, we want all reflected frequencies to become inaudible within approximately 500 ms. There is little point in covering our walls with excessive absorbent material, since absorbers are most effective at high frequencies—these are readily absorbed by normal soft materials as well.

And so, in order to minimize the effect of room modes, the key is to target the low frequencies. Both low-frequency diffusers and absorbers are an issue in small rooms since in order to be effective they have to be of considerable depth. For example, in order to absorb 85 Hz, an absorber would have to be around 1 m deep. Companies such as RPG, RealTraps, Auralex, and many others offer affordable, relatively small bass traps that fit most project studios and provide good damping of room modes. It should be pointed out that placing bass traps in a room will not reduce the bass response in any unwanted way. The first reason for this is that by minimizing the effect of standing waves, a smoother frequency response is achieved throughout the room, which in turn means that, at various points, low frequencies will be heard better. Second, bass traps help in reducing the decay time of reflected low frequencies (even non-resonant ones), which, in small, domestic rooms, is usually longer than desired.

Flutter echo

While room modes are the result of interaction between reflected low-frequency waveforms, flutter echoes are caused by mid- and high frequencies bouncing between two parallel reflective surfaces. In small rooms with reflective surfaces, we can clap our hands in order to produce such an effect. It sounds like quick, distinctive echoes with a metallic, ringing nature, similar to the jumping sound effect in cartoons.

Track 8.8: Flutter Echo
Flutter echo caused by handclaps in a domestic room.

Track 8.9: Snare No Flutter
The source snare track.

Track 8.10: Snare Flutter
An artificial simulation of flutter echo using a 30 ms delay (which roughly corresponds to walls 10 m apart) with feedback set to 50 percent. Although the effect is exaggerated in this track, it demonstrates the timbre coloration that can occur in real life.

Snare: Toontrack EZdrummer

The addition of flutter echo to percussive instruments, such as snares, is like the addition of non-musical delay, which colors the timbre and adds some metallic tail. For project studios, absorbers are the most practical solution to treat flutter echo. These are placed on the offending walls.

Early reflections

Reflections bouncing from nearby surfaces blend with the direct sound. Since they travel a longer distance, they arrive with a specific phase relationship to the direct sound. Mid and high frequencies are more problematic due to their shorter wavelengths. Early reflections happening within the first 40 ms cause comb filtering—a set of boosts and cancellations across the perceived frequency response. There is also smearing of the stereo image due to the delayed arrival of the direct sound and its reflections. Early reflections commonly bounce from the desk, sidewalls, and ceiling. But sound waves are not limited to specular travel—a sound wave hitting the edge of a computer screen can radiate back to the listening position as well.

Apart from effective positioning, which is discussed next, absorbers are used to suppress these early reflections. Since the mid and high frequencies are the most problematic ones, the absorbers do not have to be extremely deep—a 50 mm acoustic foam usually gives the desired results. Prioritizing the placement of these tiles is based on two factors: the nearest surfaces are treated first, and in the case of walls absorbent material is first placed halfway between the sound source and the listening position. When possible, nearby objects (such as computer screens) are protected by absorptive material, even if the reflections they cause are not specular.

> **Track 8.11: Left Right 1 kHz**
> Depending on the monitoring environment, this track might, or might not, demonstrate problems caused by early reflections. A 1 kHz sine toggles between the left and right channels in a 2-second interval. Ideally, the test tone should appear to come clearly from the speaker it is emitted from. Also, there should be no difference in level, tone, or stereo width and focus between the two speakers. While seated in the listening position, moving your head around might reveal differences between the two speakers.

Positioning monitors

Among all the acoustic-related enhancements we can apply, monitor positioning is one of two things that do not cost money, the other being reading the manual for our monitors. Nearly all monitor manuals include some practical advice regarding ideal placement and configuration of the speakers.

Where in the room?

Perhaps the most crucial positioning decision is where exactly the listening position is, which is largely determined by the position of the monitors themselves. Room modes affect the frequency response at the listening position and the ability of the speakers to reproduce specific frequencies. Comb filtering also has an effect on what we hear. Since the problems caused by room modes and comb filtering are more profound in small rooms, minor changes to the listening or monitor position in such rooms can have a dramatic effect. This makes the monitor and listening position even more crucial in small mixing environments, such as most project studios. Unfortunately, it is in these project studios where space limitations provide little or sometimes no option for monitor placement. Ideally, we would like to try out different listening and monitor positions in an empty room, usually while playing a familiar mix and test tones. It takes three people moving about—two holding the monitors and the other listening. Despite the time and technical issues involved in connecting the monitors and moving them around, the benefit from correct positioning could be invaluable.

Usually, there is only one dimension in question. The height at which we intend to sit when mixing will determine the height of the monitors. Since left/right symmetry results in a more accurate stereo image, the listening position is often halfway between. Unfortunately, halfway between any two walls is often where profound frequency imbalance is caused by a fundamental room mode. One thing we can do is to experiment to see whether it is the length or the width of a rectangular room that gives us fewer problems. Since the height is fixed, and on the left/right axis we usually sit in the center, all we have to experiment with is the front/back movement of our listening position and our speakers. It is suggested by some that the ideal listening position might be one-third of the way along the dimension. Nothing, however, beats experimentation.

The equilateral triangle

The most common monitor setup involves an equilateral triangle. The monitors are placed on two vertexes, facing the listener with their rays meeting at a focal point
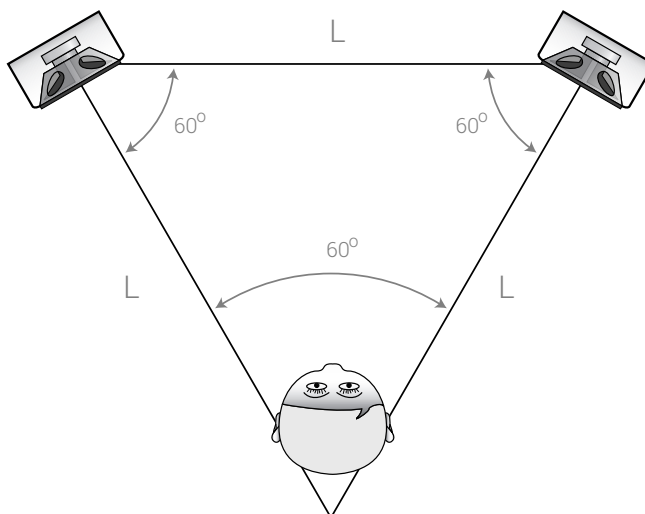
Figure 8.4 An equilateral triangle speaker setup. The angles between the speakers and the focal point behind the listener's head are all 60°. This creates an arrangement where the distance between the two speakers is equal to the distance between each speaker and the focal point.

right behind the listener's head (Figure 8.4). Many position the monitors in such a way by using a string to measure an equal distance between all vertexes (the two speakers and the focal point). It is vital that the speakers are angled toward the listener—parallel speakers produce an extremely blurred stereo image with an undefined center. Although the equilateral triangle is a recommended standard, it is worth experimenting with variations involving a wider angle at the focal point. An isosceles with 90° will result in a wider stereo image, but sometimes at the price of an unfocused center.

After the initial positioning of the speakers, it is worth playing a familiar mix and moving the head back and forth to find the best listening position. Moving the head backward will narrow the stereo image, while moving it forward will result at some point in a distorted stereo spread and blurred center. There is usually only one point where the stereo image appears to be optimal.

How far?

Once the optimal focal angle has been determined, we can move the speakers closer or farther away from the listener while sliding them on the imaginary isosceles sides. While the ear should always be the final judge, a few points are worth considering:

• It takes time for the waves from the different drivers to unite into a cohesive sound. If the speakers are too close to our ears, we can hear the sound as if coming individually from each driver (e.g., highs from the tweeter, lows from the woofer). In such cases, small head movements would result in great changes of the perceived sound.

- The closer the speakers, the more phase differences between the left and right ears, which results in a less solid stereo image.
- The farther away the speakers are, the wider the stereo image becomes, which makes panning decisions easier and reverbs somewhat more defined.
- The farther away the speakers are, the smaller the direct sound ratio is compared with the reverberant room sound. In a small room with profound resonant room modes, this is not desirable.
- Considering the wall behind the speakers, the farther the speakers are from the listener (therefore closer to the wall), the louder the low frequencies will be. This is caused by low frequencies bouncing from the back wall and returning to super-impose on the direct sound. As already mentioned, many active designs feature a bass roll-off switch to compensate for this phenomenon.
- The farther away from the listener, the closer the speakers are likely to be to the back and sidewalls, which can result in more comb filtering.

Horizontal or vertical?

Experts strongly recommend mounting monitors vertically. If the monitors are mounted horizontally, side movements of the head result in individual distance changes from each of the drivers, and thus unwanted coloration of the frequency response. If, for whatever reason, monitors are mounted horizontally, the general recommendation is for the tweeters to be on the outside. This ensures a wider stereo image, and some also claim a better bass response.

Very high frequencies are extremely directional; it can easily be demonstrated how a frequency of 18 kHz coming from a single speaker can only be heard if the ear is placed in a very specific point in space. Therefore, tweeters should be placed at ear level, or, if the monitor is angled, the tweeter's normal should be directed toward the ear.

Damping monitors

Various devices are used to separate the monitors from the surface on which they are mounted. Most widespread are isolation pads made of dense acoustic foam (Figure 8.5) and metal spikes. Monitor decouplers function in two principal ways. First, they isolate the monitor from the stand (or desk), ensuring that the speaker operates independently, with no back-vibrations from the stand interfering with the monitor operation. Second, they prevent transmission of vibrations onto the stand, which can generate unwanted resonance. The resonance of a hollow metal stand can easily be demonstrated if you clap your hands next to it. Such stands are designed to be filled with sand, which increases their mass and minimizes resonance. Also, sound generally travels faster through solid matter than through air. It is possible for sound to travel through the stand and floor and reach our body before the sound traveling through air does, possibly confusing our perception of low frequencies. Both foam and spike decouplers are fairly cheap, yet are known to have a highly positive effect on monitor performance. Mostly, they yield a tighter, more focused bass response.

Figure 8.5 The Auralex MoPAD monitor isolation pads. Both a Yamaha NS10 and a Genelec 1029 are shown here resting on the Auralex MoPADs.

A/B realm

Virtually every serious mixing studio has more than one pair of monitors. Sometimes there might even be more than one pair of near-fields. Project studio owners can also benefit from having more than one pair of monitors. Different brands vary in characteristics such as frequency response, detail accuracy, depth imaging, and stereo imaging. Having more than one brand lets us compare various aspects of the mix using different references, and can help us make important mixing decisions, especially while stabilizing the mix. Different products on the market, often called *control room matrices* or *command centers*, let us toggle between the different pairs of speakers. The Mackie Big Knob in Figure 8.6 is one of them. Cubase provides internal functionality that achieves the same, provided the audio interface has more than one stereo output.

Figure 8.6 The Mackie Big Knob. Among the features of this studio command system are monitor selection and mono summing.

Very often, when a few sets of monitors are installed, only one is used for the majority of the mixing process, with the others used for reference or specific tasks. It is important to remember that each pair of monitors will sound different. This can make frequent jumping between one set to another more confusing than helpful—as mixes sound different on each set, there's no solid reference for decision making and one might be applying changes simply because of the sonic differences between two sets. It is, therefore, advisable to take a short break before changing to a new monitoring system; this way, the sonic differences between the monitors will be less noticeable.

Headphone mixing

Having used monitors exclusively, some professional mixing engineers might be shocked to find a section on mixing using headphones in a book on mixing. Others may understand, having used headphones as a reference on occasion.

Truth be told, many people nowadays mix in bedrooms, where noise-level considerations exist (late hours, neighbors, etc.), and rooms may be poor for listening. Thus, mixing on headphones is not only sometimes the only option, but it may actually have some advantages. Some of the readers of this book, as it turns out, use headphones exclusively while mixing, so the disadvantages of doing so also need to be discussed.

Headphones vs. speakers: The theory

Before we delve into the practical side of headphone mixing, here is a short theoretical introduction.

When listening on headphones, both the left and right ears are fed exclusively with the corresponding channel. This means that the left-channel signal reaches the left ear only, and the right-channel signal reaches the right ear only (Figure 8.7a).

(a)   (b)
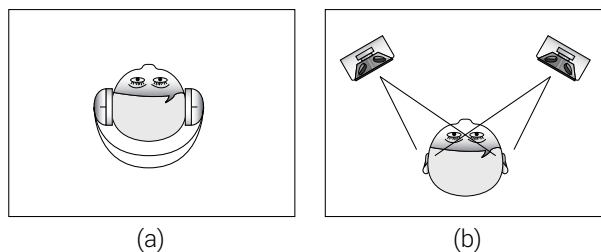
Figure 8.7 Headphones vs. speakers. With headphones, each ear is fed with the corresponding ear channel only. With speakers, each channel reaches the near ear first, and shortly after the far one.

With speakers, this is not the case. Sound from each speaker reaches the nearest ear first and the farther ear soon after (Figure 8.7b). Effectively, each ear gets the signal from the nearest speaker blended with the slightly delayed signal from the other speaker. This results in the following:

- Sounds from one speaker can mask sounds from the other.
- Overall smearing of the image for any sounds not panned to the extremes. Most severe smearing happens with sounds panned center.
- Curvature of the sound image as center-panned sounds appear deeper due to the delay between the late arrival of sound from the far speakers.

None of this happens with headphones, but stereo was conceived with speakers in mind, and for many decades now music has been made assuming playback through speakers. Our equipment, notably the design of our pan-pots (but also that of stereo effects such as reverbs), assumes the same. Mixing engineers mix using, and for, speakers. But how do these mixes translate onto headphones?

The key difference between listening to music through speakers and headphones has to do with the way our brain localizes sounds. How this happens is based on the findings of Helmut Haas and is implemented through Alan Blumlein's invention of stereo, both explored in greater detail in Chapter 14. For now, it is sufficient to say that a sound from one speaker will completely mask a similar sound from the opposite speaker if the former is approximately 15 dB louder. Practically, if the signal sent to one speaker is roughly 15 dB softer than a similar sound sent to the opposite speaker, the sound will appear to be coming entirely from the louder speaker. But with headphones, no such masking occurs, as the sound of each channel doesn't arrive at the opposite ear. To make a sound appear as if it is coming entirely from one ear, roughly 60 dB of attenuation is required on a similar sound arriving at the other ear.

In the way pan-pots are designed, when one pans from the center to the left, one should expect the location of the sound to correspond to the pot movement. This does not happen with headphones, where the sound seems to shift only slightly off the center first, before quickly "escaping" to the extreme just before the pan-pot reaches its extreme position. It is hard to place sounds anywhere between the slightly off-center position and the very extreme when listening on headphones. In fact, in the way standard pan-pots work, it is next to impossible. Positioning instruments on the sound stage is much easier when listening through speakers. Applications such as Logic offer "binaural" pan-pots, which tackle this exact problem and can achieve much better localization using headphones, but the penalty is that they do so by altering the frequency content of the signal sent to each ear, thereby altering the overall frequency spectrum of the instrument. Also, these types of "binaural" mixes do not translate well on speakers. In addition to all that, the sound stage created by speakers is limited to the typical 60° between the listener and the two speakers. With headphones, the sound stage spans 180°.

Mixing engineers work hard to create sound stages in mixes using speakers. When these mixes are played through headphones, these sound stages appear completely distorted. While this does not seem to bother most listeners, most serious music buffs insist that listening to music via speakers is far more pleasing, largely due to the lack of spatial sense when using headphones.

The dominance of speaker mixes was never questioned until not long ago, when portable MP3 players and their integration with cellular phones became so widespread. It is a valid question to ask why we still mix using (and for) speakers when so many

people nowadays listen via headphones. There is an unexploited opportunity here for record labels to produce "speaker" and "headphone" versions. This would make sense not only from a mixing point of view, but also from mastering, consumer, and label revenue points of view.

The advantages

Some recording and mixing engineers take their reference headphones to studios they are not familiar with. Headphones provide a solid reference, and their sound only alters when different headphone amplifiers are used. As previously explained, the room plays a dominant part in what we hear with a speaker setup—the sound headphones produce is not affected by room acoustics or modes.

This is very important for rooms with flawed acoustics, such as many bedrooms and project studios. In such rooms, a good pair of headphones, together with a good headphone amp, can be a real aid. Having room modes out of play means that the range between 80 and 500 Hz can be more evenly reproduced, although studio monitors still have an advantage reproducing very low frequencies compared with most headphones. Other acoustic issues simply don't exist when using headphones; for example, the comb filtering caused by early reflections, the masking between the left and right speakers, and even the directivity of the tweeters. It can be generalized that, as far as frequency reproduction is concerned, the results we get using good headphones are more accurate than those generated by speakers in a flawed room.

For many compression, gating, and dynamic range tasks, speakers do not provide a clear advantage over headphones—as long as the processing does not affect the perceived depth of the instrument, headphones can be useful.

The disadvantages

While headphones can be great when dealing with frequencies and useful when treating certain dynamic aspects of a mix, there are also a few disadvantages in using them, and they are almost useless for some mixing tasks.

As discussed, the spatial image created by headphones is greatly distorted and conventional tools make it very hard to craft appropriate sound stages on headphones. Any sound stage decisions we make, whether left/right or front/back, are better made using speakers. As depth is often generated using reverbs, delays, or other time-based effects, the configuration of these effects benefits from using speakers.

Earlier in the chapter, it was stated that ideal mixing rooms aim to have the reverb within them decaying at around 500 ms and that anechoic chambers, where no reverb exists, are far from ideal. As we lose the room response when using headphones, we also lose the reverb, so it is as if we mix in a space similar to an anechoic chamber. Most people find this less pleasant. The lack of room reverb, together with the close proximity of the headphones' diaphragms to our eardrums, mean that ear fatigue is more likely to occur. At loud levels, headphones are also more likely to cause pain and damage to the eardrum.

Headphone amps

In the same way as power amps influence the performance of passive speakers, headphone amps influence the performance of headphones. Cheap headphone amps colorize the sound, can cause audible distortion (especially at loud levels), and tend to present an undesired correlation between frequency response and listening levels. A good set of headphones together with a good headphone amp can easily cost as much as a good, active near-field monitor.

Which headphones?

Earphones are small constructions with very small diaphragms that render them unsuitable for critical listening tasks. Thus, this section focuses on headphones, for which there are three principal designs:

- **Closed-back**—this design aims to achieve maximum isolation between the sound generated by the headphones and the environment. This is useful in recording situations when a minimum amount of headphone leakage is sought (vocal overdubs, for example). It is also useful when a bassist stands next to the drummer in a studio where the loud drums mask the sound produced by the headphones. Outside studios, this design is also used in noisy environments. In order to achieve maximum isolation, the cups have their back sealed and the headband normally presses the cups against the head. The closed cup provides some design challenges as it means some sound reflects back to cause comb filtering. This design is also normally unsuitable for long listening periods as the pressure applied to the head can cause discomfort.
- **Open-back**—this design puts audio quality and comfort first, making it most suitable for critical listening applications and the long listening periods that mixing invariably requires. The back is open to prevent comb filtering and so the sound normally travels both ways easily with little isolation. For maximum comfort, most reference headphones will have cushions bigger than the pinna (the outer, visible part of the ear), and the headband normally applies little pressure to the head around the ears. Long-term usage can result in pain at the top of the head against which the weight of headphones applies, so more expensive designs are lighter. Similarly, cushions on better designs are also made from materials that absorb sweat yet dissipate heat for additional comfort.
- **Semi-open**—this is a marriage between the two previous designs, aimed at providing moderate isolation when such a thing is needed. An example is live sound, when an engineer might wish to hear the sound coming from the front of house yet still be able to hear either soloed tracks or a communication feed.

9 Meters

A common saying in mixing is "Listen, don't look." Being a sonic art, mixing is all about listening. Still, meters are always in sight in mixing environments, and for good reason. There are various stages and situations where meters can be useful. Of the many types of meters in audio engineering, mixing makes notable use of two: the peak and VU meters. We will also discuss phase meters briefly. It would be hard to discuss metering without giving a little bit of background. The short section below will also be handy in later chapters.

Amplitude vs. level

In general acoustic terms, amplitude describes the extent of air pressure deviation from the normal atmospheric pressure. A microphone converts changes in air pressure to voltages. An A/D converter converts these voltages into discrete numbers. Changes in air pressure happen above and below the normal atmospheric pressure (the zero reference), resulting in sound amplitude that is *bipolar*—it has both positive and negative magnitudes. The voltages and numbers used to describe audio signals are also bipolar. An audio system has an equal positive and negative capacity. Professional audio gear, for example, uses –1.23 to +1.23 volts; an audio sequencer uses the –1 to +1 rational range.

One of our main interests is to make sure a signal will not exceed the limits of a system. For example, we do not want the samples within our audio sequencer to rise above +1 or drop below –1. But we do not really care whether the signal exceeds the positive or negative limits; we only care about its *absolute* magnitude. The term *level* in this book denotes the absolute magnitude of signals. Amplitudes of +0.5 and –0.5 both denote a level of 0.5. Figure 9.1 demonstrates these differences.

Using numbers and voltages to express signal levels would be hugely inconvenient. A professional desk indicating that the signal level is 0.307073056039340737757909071411 V, or an audio sequencer showing us that the sample value is 0.25, is rather overwhelming, especially considering the two denote exactly the same level, only in different units. The decibel system provides an elegant solution—it expresses levels with more friendly values that mean the same on all systems. The two numbers above are simply –12 dB.

The limit, or the standard operating level of a system, is denoted by 0 dBr (dB reference). On professional audio equipment, 0 dBr is a level of 1.23 V; within an audio sequencer, it is a
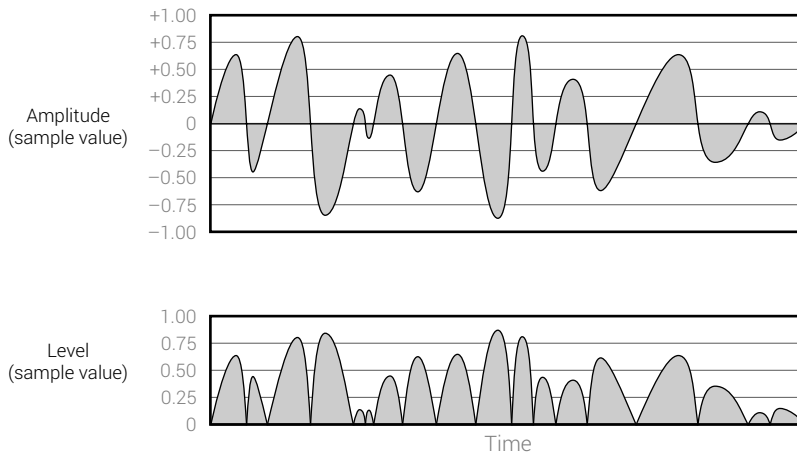
Figure 9.1 Amplitude and level. The top graph shows the amplitude of the waveform, which has both positive and negative magnitudes. The level representation is simply the absolute magnitude of the signal.

sample level of 1. Since 0 dBr is the highest limit of the system, levels are mostly negative. But on analog equipment, signals can go above 0 dBr—they might clip, they might distort, but they can still go there—so we also have positive levels. We call the range above 0 dBr *headroom*. Even within an audio sequencer, signals can go above 0 dBr, but for a few good reasons we are made to believe that they cannot (more on this in Chapter 11). For now, we should regard 0 dBr as the absolute limit of a digital system.

For convenience, the "r" from dBr is omitted henceforth.

Mechanical and bar meters

Mechanical meters involve a magnet and a coil that move a needle. They take up quite a lot of space, and usually only involve a scale of around 24 dB. Bar meters involve either a column of LEDs, a plasma screen, or a control on a computer screen. Bar meters might provide some extra indicators in addition to the standard level gauge (Figure 9.2):

- **Peak hold**—a held line on the meter indicating the highest meter reading. Usually, the hold duration can be set to forever, a few seconds, or off. This feature tells us how far below 0 dB the highest peak of the signal is, which can be useful during recording or bouncing, where we use this information to push levels farther up.
- **Peak level**—a numeral display that shows the level of the highest peak.
- **Clip indicator**—an indicator that lights when the signal exceeds the clipping level, which is normally set to 0 dB on a digital system. On analog equipment, the clipping level might be set above 0 dB, to the level where the signal is expected to distort.

Clear hold and *clear clips* functions are available for all three facilities.

Clip indicator
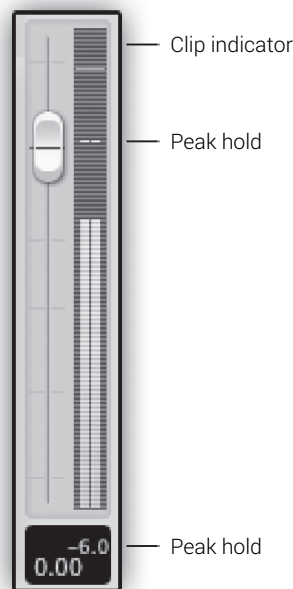
Peak hold

Peak hold

Figure 9.2 Cubase track meter facility. In addition to the standard level meter, there is also a graphical and numeral peak hold indicator and a clip indicator.

Peak meters

Peak meters are straightforward—they display the instantaneous level of the signal. Their response to level changes is immediate. Peak meters are mandatory when the signal level must not exceed a predefined limit, such as 0 dB on a digital system. In essence, this is their main role.

On a digital system, the highest level of a peak meter scale is 0 dB. As the lowest level may vary with relation to the bit depth (being –96 dB for 16-bit audio, –144 dB for 24-bit audio, and so on), many digital meters have their bottom level set around –70 dB. If a peak meter is installed on analog equipment, the scale might extend above 0 dB.

Average meters

Our ears perceive loudness in relation to the **average level of sounds, not their peak level** (Figure 9.3 demonstrates the difference between the two). One disavantage of peak meters is that they teach us little about the *loudness* of signals. In order to bear any resemblance to loudness, a meter has to incorporate some averaging mechanism. There are various ways to achieve this. A mechanical meter might employ an RC (resistor-capacitor) circuit to slow down the rise and fall of the needle. In the digital domain, the mathematical root mean square (RMS) function might be used. Regardless, the movement of the meter should roughly reflect the loudness we perceive. One example where this can be useful is when compressing vocals—it is the perceived loudness of the vocals we want even.
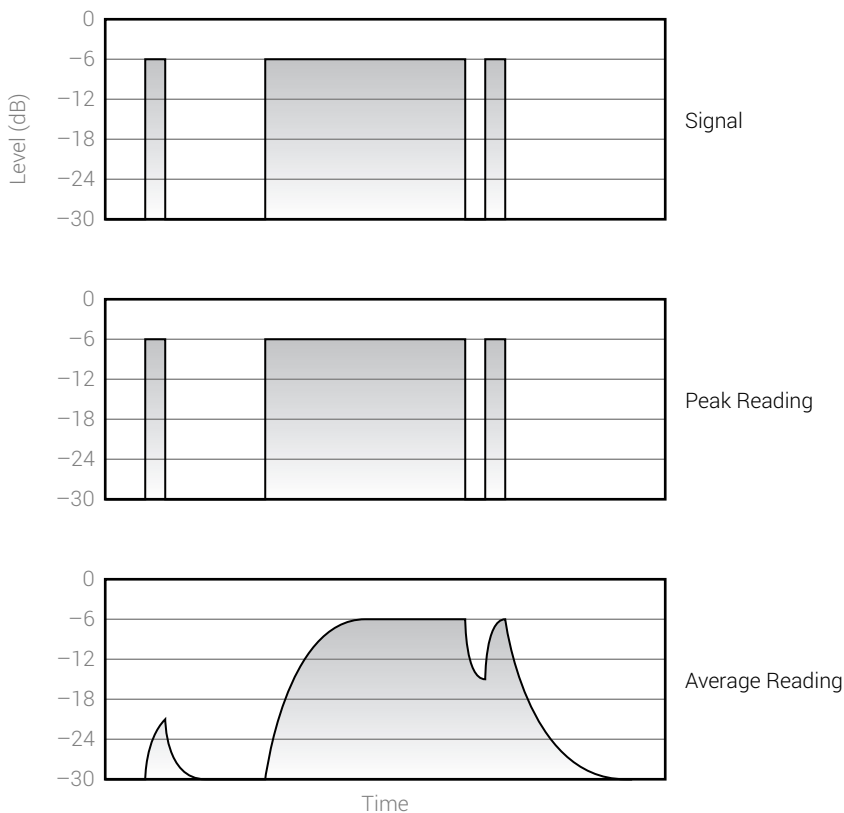
Figure 9.3 Peak vs. average readings. A peak meter tracks the level of the signal instantaneously; therefore, its reading is identical to the signal level. An average reading takes time to rise and fall, but the resultant readout is reflective of our perception of loudness. We can say that an average readout seems lazy compared to a peak readout.

The reading of averaging meters relates to perceived loudness.

Of all the averaging meters in mixing, the mechanical VU (volume unit) meter is the most popular. Figure 9.4 shows a plugin version. The scale spans –20 to +3 dB, and it is worth noting the distribution of different levels. For example, the right side of the scale covers nearly 6 dB, while the left covers the remaining 17 dB. Many studio engineers are accustomed to this type of meter, and some even have enough experience to set rough levels of various instruments just by looking at them.

VU readings can also be displayed on bar meters. This simply involves showing the RMS reading on the standard peak scale (–70 to 0 dB). However, since people are so used to the –20 to +3 range of the mechanical VU meters, a VU bar meter might have its scale covering these 23 dB only (Figure 9.5).

Figure 9.4 The PSP *VintageMeter*. This free plugin offers averaging metering characteristics similar to those on the mechanical VU meters.

Figure 9.5 The Sonalksis *FreeG* plugin. This free plugin provides extended metering capabilities, among other features. The main meter characteristic is based on one of four ballistics settings (seen on the right panel); these include types of meters not discussed here, such as the broadcast BBC type. The narrow bar within the main meter shows RMS levels. There are numeral and graphical (arrows) peak hold indicators for both peak and RMS.

Phase meters

VU or peak meters are often provided per channel and for the stereo mix. One more type of meter worth briefly mentioning is the phase meter. Phase meters are a
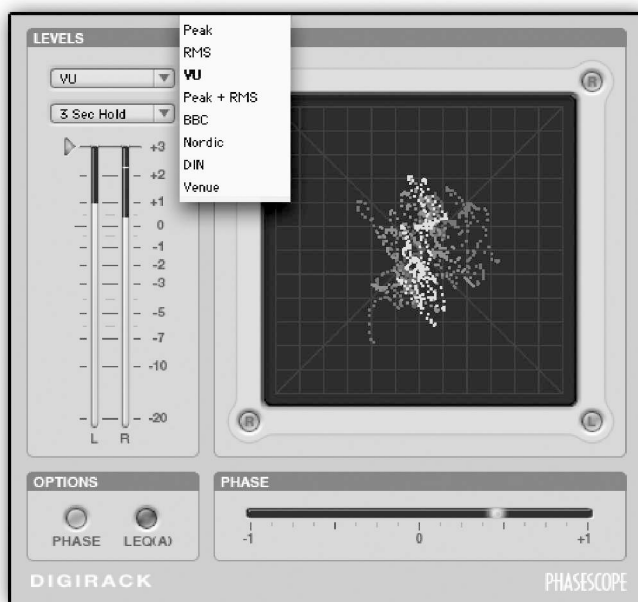
Figure 9.6 The Digidesign *PhaseScope* plugin. The phase meter can be seen at the bottom. Also note the –20 to +3 dB VU scale on the meters, and the number of possible meter characteristics. The psychedelic graph is actually a very useful tool called the Lissajous curve. It provides visual imaging of the stereo signal, from which we can learn about stereo image shifting, stereo width, phase problems, and more.

common part of large-format consoles. They meter the phase relationship between the left and right channels of the mix. The meter scale ranges from –1 to +1. The +1 position denotes that both the channels are perfectly in phase; that is, that they output exactly the same signal (which is essentially mono). Zero denotes that each channel plays something completely different (essentially perfect stereo) and –1 denotes that the two channels are perfectly phase-inverted. Generally speaking, positive readings tell us that our mix is phase-healthy, whereas negative readings suggest that there might be a problem. We want the meter to remain on the positive side of the scale unless we deliberately use an effect that involves phase inversion between the left and the right speakers (like the out-of-speakers effect that is described later in Chapter 12) (Figure 9.6).

10 Mixing consoles

This chapter covers the common mixing (not recording) functionality offered by a typical console. It provides vital background information for Chapter 11, on software mixers, and the chapters succeeding it.

A *mixing console* (or a *desk*—both terms are interchangeable) is an independent hardware device used alongside a multitrack recorder at the heart of any recording or mixing session in professional studios. Consoles vary in design, features, and implementation. Products range from compact 8-channel desks through 96-channel large-format analog consoles to large-format digital consoles that can handle more than 500 input signals. In a mixing session, the console's individual channels are fed with the individual tracks from the multitrack recorder. (It is worth noting the terminology—a *channel* exists on a console; a *track* exists on a multitrack recorder.) The mixing console offers three main functionalities:

- **Summing**—combining the audio signals is the heart of the mixing process; most importantly, various channels are summed to stereo via the mix bus.
- **Processing**—most consoles have on-board equalizers, while large-format consoles also offer dynamic processors.
- **Routing**—to enable the use of external processors, effects, and grouping, consoles offer routing functionality in the form of insert points, auxiliary sends, and routing matrices.

All mixing consoles have two distinctive sections:

- **Channel section**—a collection of channels organized in physical strips. Each channel corresponds to an individual track on the multitrack recorder. The majority of channels on a typical console support a **mono input**. In many cases, all the channels are identical (functionality and layout), yet on some consoles there might be two or more types of channel strips (for example, mono/stereo channels, channels with different input capabilities, or channels with different equalizers).
- **Master section**—responsible for central control over the console, and global functionality. For instance, master aux sends, effect returns, control room level, and so on. Large-format analog consoles might also have a **computer section** that provides automation and recall facilities.

Buses

A *bus* (or *buss* in the United Kingdom) is a common signal path to which many signals can be mixed. By way of analogy, a bus is like a highway into which many small roads flow. A summing amplifier is an electronic device used to combine (mix) the multiple sources into one signal, which then becomes the bus signal. Most buses are either mono or stereo, but surround consoles also support multichannel buses; for example, a 7.1 mix bus. Typical buses on a large-format console are:

- mix bus;
- group buses (or a single record bus on compact desks);
- auxiliary buses; and
- solo bus.

Processors vs. effects

The devices used to treat audio signals fall into two categories: processors and effects. It is important to understand the differences between the two.

Processors

A processor is a device, electronic circuit, or software code used to alter an input signal and *replace* it with the processed, output signal. Processors are used when the original signal is not required after treatment. Processors are therefore connected in series to the signal path (Figure 10.1). If we take an equalizer, for example, it would make no sense to keep the original signal after it has been equalized. Examples of processors include:

- equalizers;
- dynamic range processors;
- compressors;
- limiters;
- gates;
- expanders;
- duckers;
- distortions;
- pitch correctors;
- faders; and
- pan-pots.
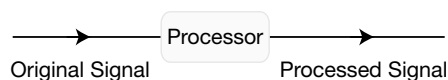
Figure 10.1 Processors are connected in series to the signal path and replace the original signal with the processed one.

Effects

Effects **add** something to the original sound. Effects take an input signal and generate a new signal based on that original input. There are then different ways to mix the effect output with the original one. Effects are traditionally connected in parallel with the signal path (Figure 10.2). Most effects are time-based, but pitch-related devices can also fall into this category. Examples of effects include:

- Time-based effects:
  - reverb;
  - delay;
  - chorus; and
  - flanger.
- Pitch-related effects: pitch shifter and harmonizer.
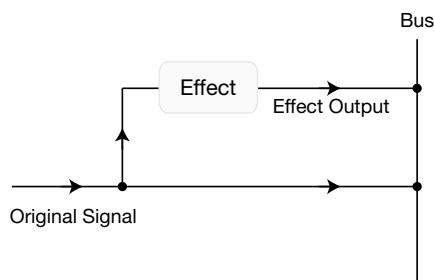
Figure 10.2 In standard operation, a copy of the original signal is sent to an effects unit. The original signal and the effect output are then often mixed when summed to the same bus.

Within an effects unit, we distinguish between two types of signal. The **dry** signal is simply the unaffected, original input signal. The **wet** signal is the new signal that the effects unit produces. If we take a delay, for example, the vocal sent to the delay unit makes the dry signal, and the delays generated by the unit constitute the wet signal. While logic has it that effects units should only output the wet signal (and indeed this should be the case when effects are connected in parallel), many effects let us mix internally the dry and wet signals, and output both. The mix between the dry and the wet signals is either determined by a single, percentage-based control, or two separate level controls, one for dry and another for wet (Figure 10.3).
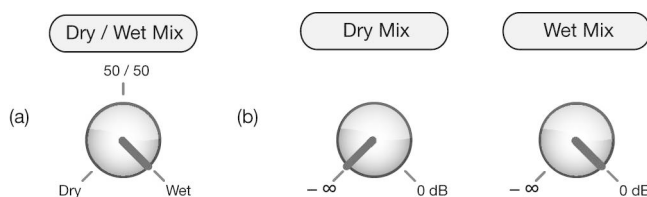
Figure 10.3 Two ways to mix the dry and wet signals within an effect box. (a) A single rotary control with fully dry signal at one extreme, fully wet signal on the other, and equal mix of dry and wet in the center. (b) Two independent controls; each determines the level of its respective signal.

Connecting processors and effects

The standard method for connecting processors is to use an **insert point**, while effects are normally connected using an **auxiliary send**. Both are described in detail in this chapter. If processors and effects are not connected as intended, some problems can arise. Still, exceptions always exist; for example, compressors might be connected in parallel as part of the parallel compression technique. In some situations, it makes sense to connect reverbs using inserts. These exceptions are explored in later chapters.

Basic signal flow

In this section, we will use an imaginary six-channel console to demonstrate how audio signals flow within a typical console, and key concepts in signal routing. We will build the console step by step, adding new features each time. Three types of illustrations will be presented in this section: a signal flow diagram, the physical layout of the console, and its rear panel (horizontally mirrored for clarity).

The signal flow diagrams introduced here are schematic drawings that help us understand how signals flow, how these can be routed, and where along the signal path key controls function. By convention, the signal flows from left to right, unless indicated otherwise by arrows. A signal flow diagram is a key part of virtually every console manual. Unfortunately, there isn't a standard notation for these diagrams, so various manufacturers use different symbols. Common to all these diagrams is that identical sections are shown only once. Therefore, a signal flow diagram of a console with 24 identical channels will only show one channel and this channel is often framed to indicate this. Most master section facilities are unique, and any repetitions are noted (for example, if groups 3–4 function exactly like groups 1–2, the diagram will say "groups 3–4 identical" next to groups 1–2). The signal flow diagrams in this chapter are simplified. For example, they do not include the components converting between balanced and unbalanced.

Step 1: Faders, pan-pots, and cut switch

The first step involves the basic console shown in Figure 10.4. Each channel is fed from a respective track on the multitrack recorder. As can be seen in Figure 10.4b, each channel has a fader, a pan-pot, and a cut switch. The signal flow diagram in Figure 10.4a shows us that the audio signal travels from the line input socket through the cut switch, the fader, and then the pan-pot. Each pan-pot takes in a mono signal and outputs a stereo signal. The stereo signal is then summed to the mix bus (the summing amplifier is omitted from the illustration) and a single fader alters the overall level of the stereo bus signal. Finally, the mix-bus signal travels to a pair of mono outputs, which reside at the rear panel of the console (Figure 10.4c).

Step 2: Line gains, phase-invert, and clip indicators

Many consoles have a line-gain pot, phase-invert switch, and a clip indicator per channel strip. In larger consoles, both the line-gain and phase-invert controls reside in a dedicated input section, which also hosts recording-related controls such as microphone-input selection, phantom power, and a pad switch. Figure 10.5 shows the addition of these controls to our console.
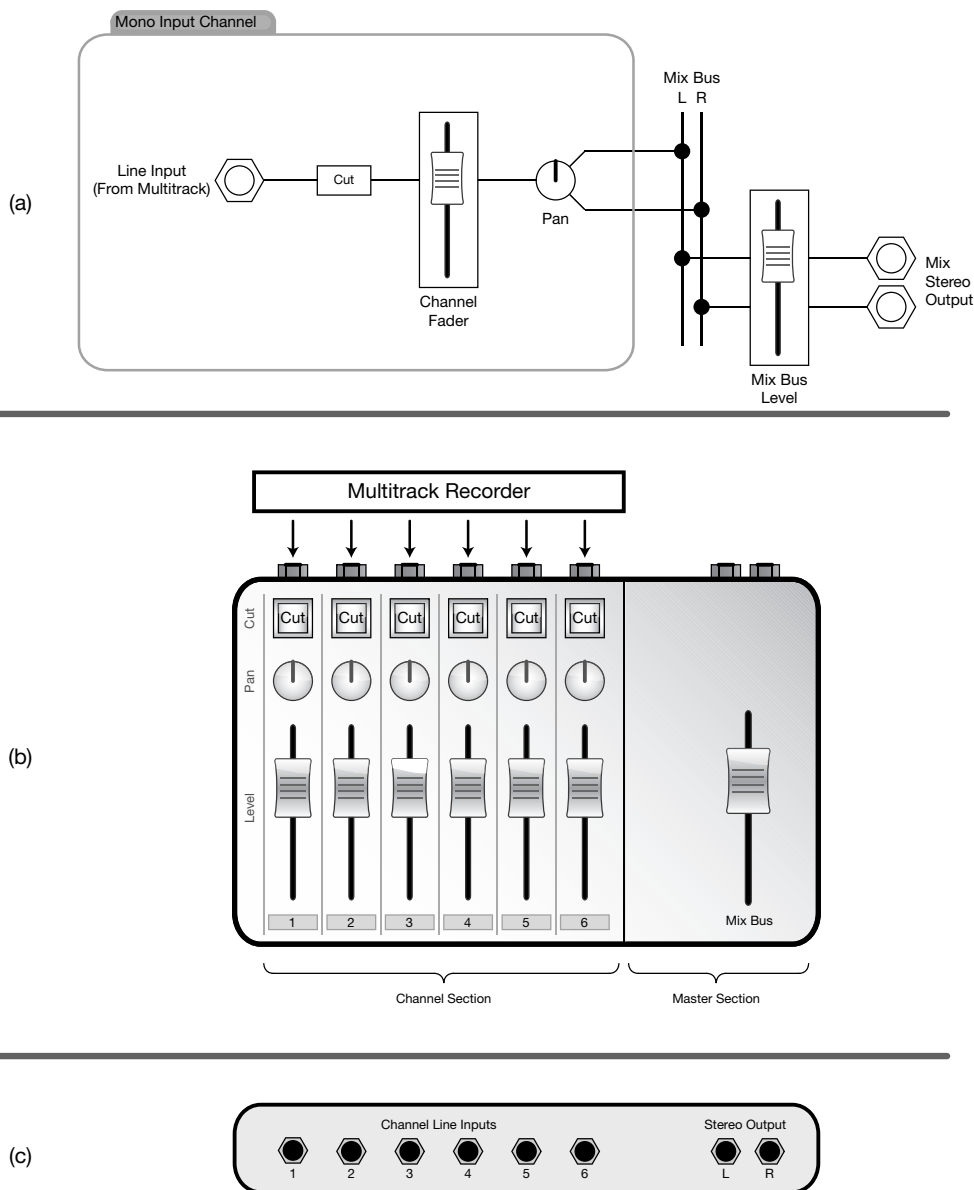
Figure 10.4 The first step in our imaginary six-channel mixer only involves a fader, pan-pot, and cut switch per channel. Note that the only control to reside in the master section is the mix-bus level.

The **line-gain** control (or *tape-trim*) lets us boost or attenuate the level of the audio signal before it enters the channel signal path. It serves the mixing engineer in two principal ways. First, it lets us **optimize** the level of the incoming signal. A good recording engineer will make sure that all the tracks on the multitrack recorder are at the optimum level; that is, the highest level possible without clipping (mostly for digital media) or the addition of unwanted distortion (for analog media). If a track level is not optimal, it would be wise to alter it at this early stage. For instance, if the recording level of the vocalist is very low, even the maximum 10 dB extra gain standard faders provide
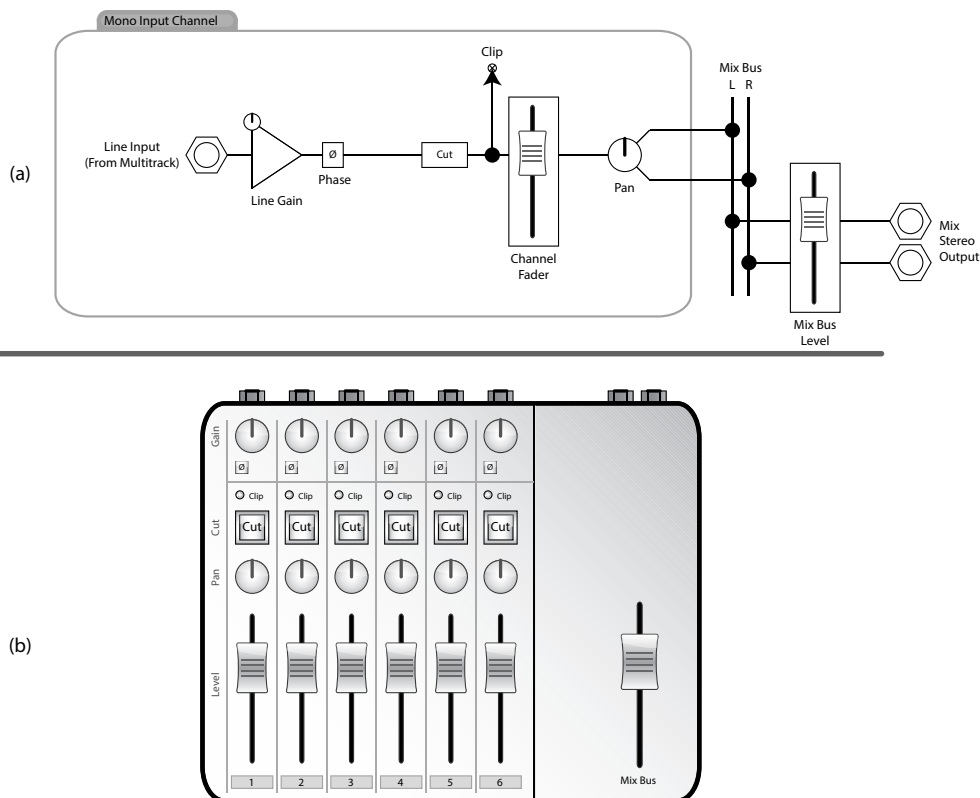
Figure 10.5 In step 2, a line-gain control, a phase-invert switch, and a clip indicator have been added per channel strip. The rear panel of the console is the same as in the previous step and not shown.

might not be sufficient to make the vocals loud enough in the mix. In such cases, we will need to bring down the rest of the mix, which entails adjusting all other faders. In rare cases, the input signal can be too hot, and might overload the internal channel circuitry. The line-gain pot is used in these situations to bring down the level of the input signal.

Funnily enough, the over-hot input signal is sought after in many situations. An old trick of mixing engineers is to boost the input signal so it intentionally overloads the internal channel circuitry, the reason being that the overloading of the internal analog components is known to produce a **distortion** that adds harmonics that are quite appealing. Needless to say, it has to be done in the right amounts and on the right instrument. Essentially, the line gain in analog consoles adds distortion capabilities per channel. One thing worth keeping in mind is that dynamic range processors, as well as effect sends, succeed the line-gain stage. Therefore, any line-gain alterations would also require respective adjustment to, for example, the compressor's threshold or the aux send level. It is therefore wise to set this control during the early stages of the mix, before it can affect other processors or effects in the signal path.

The line-gain controls found on analog consoles are used to add distortion.

A good recording engineer will also ensure that all of the tracks on the multitrack recorder are phase-coherent. In cases where a track was mistakenly recorded with inverted phase, the **phase-invert** control will make it in-phase.

The **clip indicators** on an analog console light up when the signal level overshoots a certain threshold set by the manufacturer. There can be various points along the signal path where the signal is tested. The signal flow in Figure 10.5a tells us that in our console, there is only one point before the channel fader. A lit clip indicator on an analog console does not necessarily suggest a sonic impairment, especially if we remember the over-hot signals. If we take some SSL desks, for example, it would not be considered hazardous to have half of the channels clipping when working on a contemporary, powerful mix. The ear should be the sole judge as to whether or not a clipping signal should be dealt with.

Step 3: Onboard processors

Most desks offer some onboard processors per channel, from basic tone controls on a compact desk to a filter section, fully featured four-band equalizer, and a dynamics section on large-format consoles. The quality and the number of these onboard processors dictate much of the console's value, while playing a major role in the console's sound. Figure 10.6b shows the addition of a high-pass filter, an equalizer, and a basic compressor to each channel on our console. As can be seen in Figure 10.6a, these processors are located in the signal path between the phase and cut stages, as they would appear on a typical console. Some consoles offer dynamic switching of the various processors. For example, a button would switch the dynamics section to before the equalizer.

Step 4: Insert points

The onboard processors on a large-format console are known for their high quality, yet in some situations we prefer to use an external unit. Sometimes it is simply because a specific processor is not available onboard (duckers, for example, are not very common, even on large-format consoles), and sometimes the sound of an external device is favored over an onboard processor.

As the name suggests, insertion points let us insert an external device into the signal path. They do so by breaking the signal path, diverting the existing signal to an external device, and routing the returned, externally processed signal so it **replaces** the original signal. In practice, most consoles always send a copy of the signal to the insert send socket, and the insert button simply determines whether the returned signal is used to replace the existing signal (Figure 10.7). When insert sends are not used to connect an external device, they are used as a place from which a copy of the signal is taken. For example, when adding sub-bass to a kick, a copy of the kick is needed in order to trigger the opening and closing of the gate (more on this in Chapter 19).

It is important to remember that each external unit can only be connected to one specific channel—when using the channel's insert point, it is impossible to send any other channel to the same external unit. However, it is possible to route the output of one external unit into another external unit and return the output of the latter into the console. This way, we can daisy chain external processors, for example in cases where we want both an external equalizer and a compressor to process the vocal.

(a)

(b)
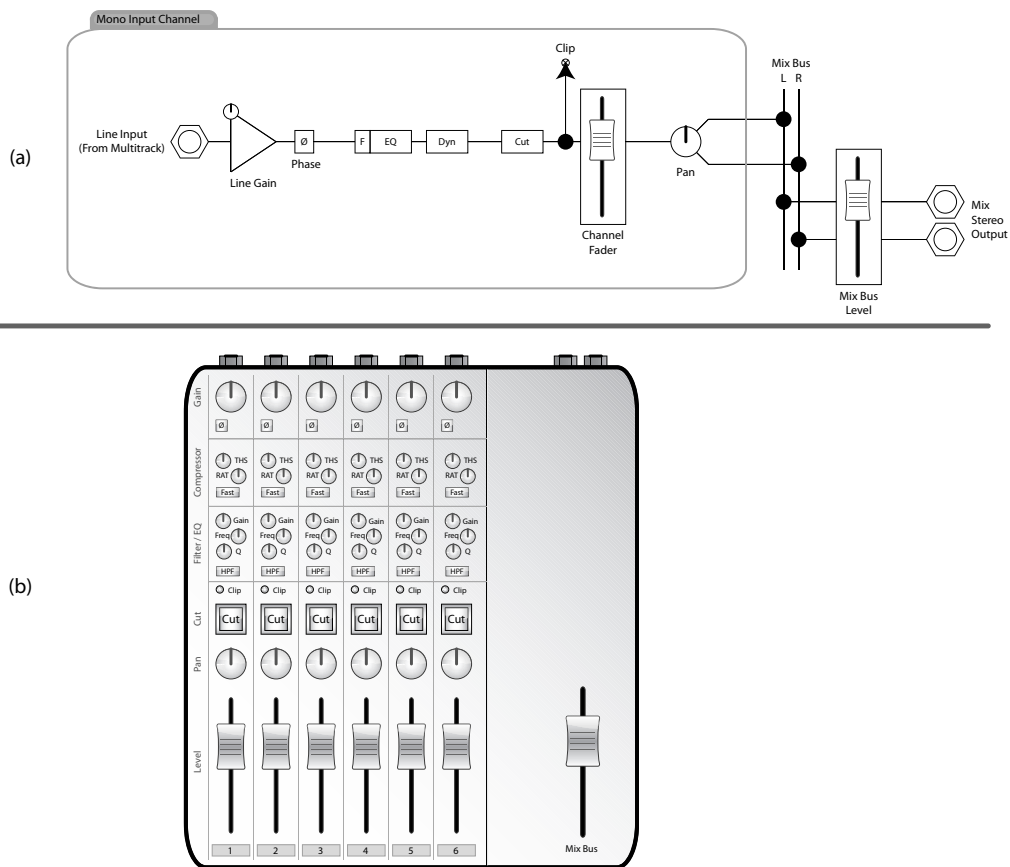
Figure 10.6 As can be seen from the desk layout (b), the additions in step 3 involve a high-pass filter, a single-band fully parametric equalizer, and a basic compressor with threshold, release, and fast attack controls.
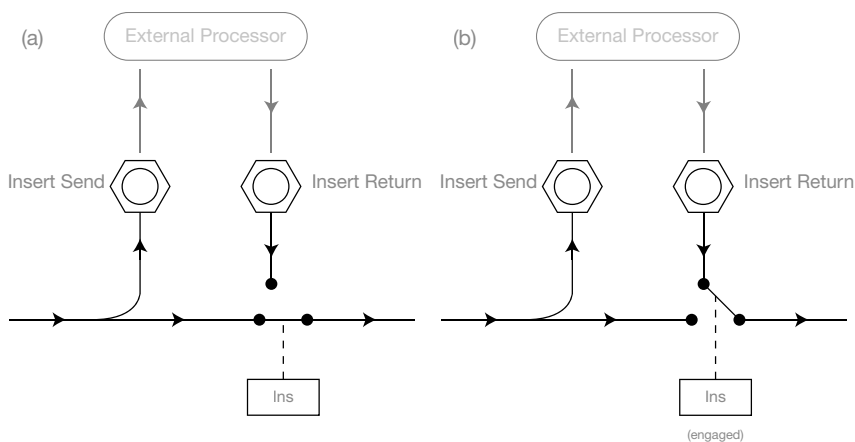
Figure 10.7 A typical insert point. (a) When the insert switch is not engaged, a copy of the signal is still routed to the insert send socket, but it is the original signal that keeps flowing in the signal path. (b) When the insert switch is engaged, the original signal is cut and the insert return signal replaces it.

Figure 10.8a shows the addition of an insert point into our console's signal flow. Note that the insert point is located after the processors and before the cut stage. Some consoles enable dynamic switching of the insert point, very often before or after the onboard processors. The only addition on the layout itself (Figure 10.8b) is the insert-enable button. On the console's rear panel in Figure 10.8c, both insert send and return sockets have been added per channel. Two balanced sockets are more common on larger consoles, but some smaller desks utilize one TRS socket that combines both the send (tip) and the return (ring). A Y-lead cable is used in this type of unbalanced connection to connect the external unit.
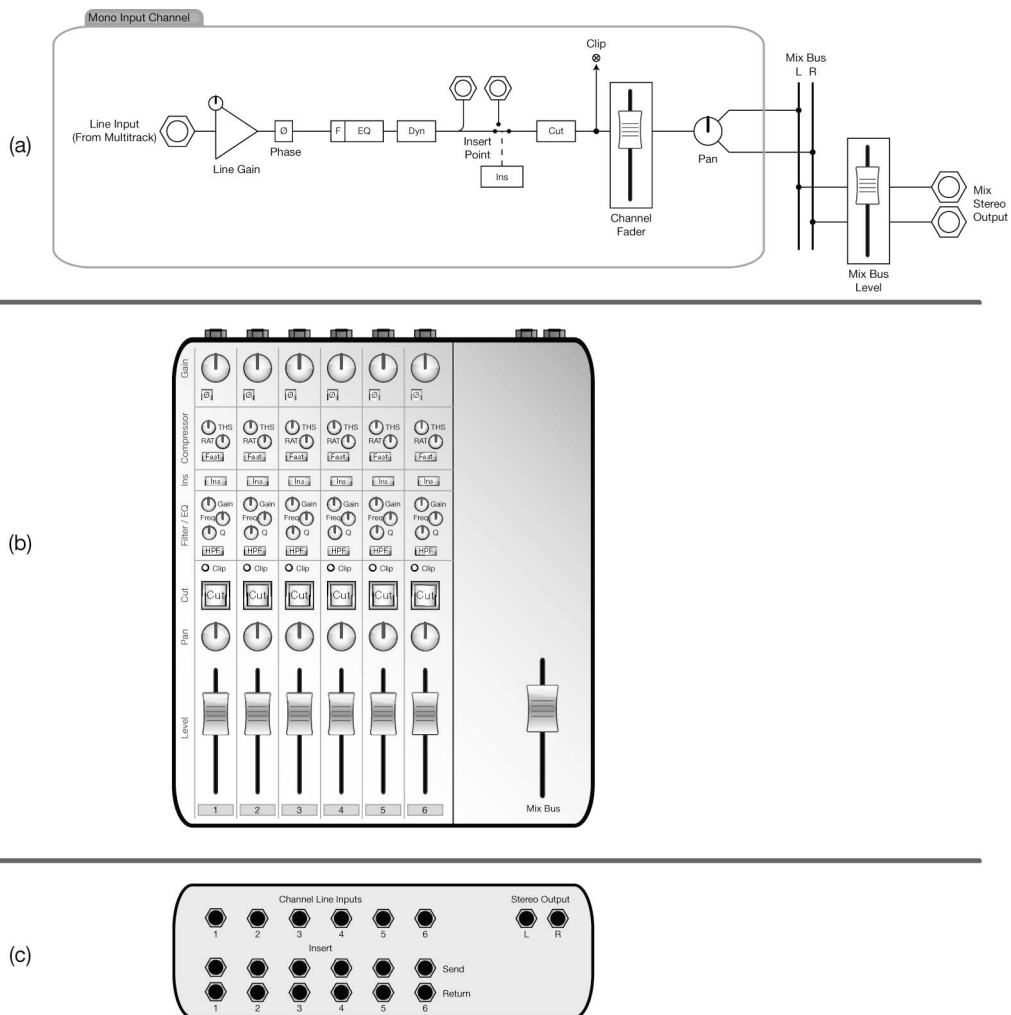
Figure 10.8 In step 4, insert points have been added to the console per channel. The only addition to the layout (b) is the insert switch, while each channel has additional send and return sockets on the rear panel.

The importance of signal flow diagrams

The console we have designed so far is very typical—both the order of stages in the signal flow and the layout of the controls on the surface are very similar to those of an ordinary console. One thing that stands out immediately is that the controls on the console's surface are organized for maximum user convenience, and not in the order of processing. Figure 10.9 illustrates how the signal actually travels if we compare it to the channel layout from step 4. We cannot tell just by looking at a desk whether the dynamic processors come before or after the equalizer, tell whether the insert point comes before or after the processors, or gather other signal-flow-related information that can be vital during the mixing process. In many cases, the manual will not answer these questions, leaving the signal flow diagram as our only reference.
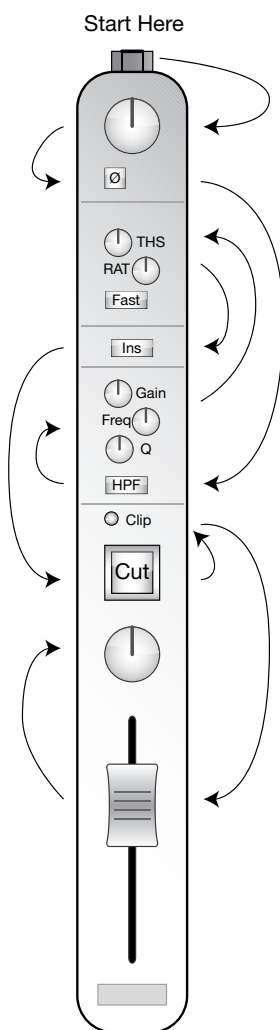
Figure 10.9 This somewhat overwhelming illustration demonstrates how the audio signal really flows through the channel controls.

Step 5: Auxiliary sends

An auxiliary send (*aux send* or *send*) takes a copy of the signal in the channel path and routes it to an internal auxiliary bus. The auxiliary bus is then routed to an output socket, which could then feed an external effects unit. Each channel strip on the console has a set of aux send-related controls. To distinguish these controls from identical auxiliary controls on the master section, we call them *local auxiliary controls*; these may consist of:

- **Level control**—there is always a pot to determine the level of the copy sent to the auxiliary bus.
- **Pre-/post-fader switch**—a switch to determine whether the signal is taken before or after the channel fader. Clearly, if the signal is taken post-fader, moving the channel fader would alter the level of the sent signal. When connecting effects such as reverbs, we often want the effect level to correspond to the instrument level. Post-fader feed enables this. When the signal is taken pre-fader, its level is independent of the channel fader and will be sent even if the fader is fully down. This is desirable during recording, where the sends determine the level of each instrument in the cue mix—we do not want changes to the monitor mix to affect the artists' headphone mix. If no switch is provided, the feed can be either pre- or post-fader.
- **Pan control**—auxiliary buses can be either mono or stereo. When a stereo auxiliary is used, there will be a pan-pot to determine how the mono channel signal is panned to the auxiliary bus.
- **On/off switch**—enables/disables the send to the auxiliary bus. Sometimes simply labeled "mute."

It is important to understand that we can send each channel to the same auxiliary bus while setting the level of each channel individually using the local aux level (Figure 10.10). For instance, we can send the vocal, acoustic guitar, and tambourine to a single auxiliary bus, which is routed to a reverb unit. With the local aux level controls, we set the level balance between the three instruments, which will eventually result in different amounts of reverb for each instrument. As opposed to insert sends, aux sends allow us to share effects between different channels.

Also, in Figure 10.10, it is worth noting the controls used to set the overall level of the auxiliary bus or cut it altogether. These **master aux controls** reside in the master section. Master aux controls are often identical to the local ones, with one exception: there is never a pre-/post-fader control. These are done on a per-channel basis and not globally for each auxiliary bus.

Most consoles have more than one auxiliary bus, and thus more than one set of local and master aux controls. On some consoles, the number of auxiliaries determines how many effects units can be used in the mix. For example, in a console with only two auxiliaries, aux 1 might be routed to a reverb unit, with aux 2 to a delay unit. As with inserts, we can daisy-chain units in order; for example, to have aux 1 sent to a delay, which is then followed by a reverb. The different aux sends are not always identical, and each might involve a slightly different set of controls. In the case of our console, two aux sends have been added (Figure 10.11). Aux 1 is stereo, has a pre-/post-switch, pan control, and a cut switch. Aux 2 is mono and only

features level and cut controls. As can be seen in Figure 10.11a, aux 1 has a fixed post-fader feed. To the desk layout in Figure 10.11b, the respective controls for each local aux have been added per channel. Also, the master section now offers both level and cut controls for each master aux. Figure 10.11c shows the addition of output sockets per auxiliary.
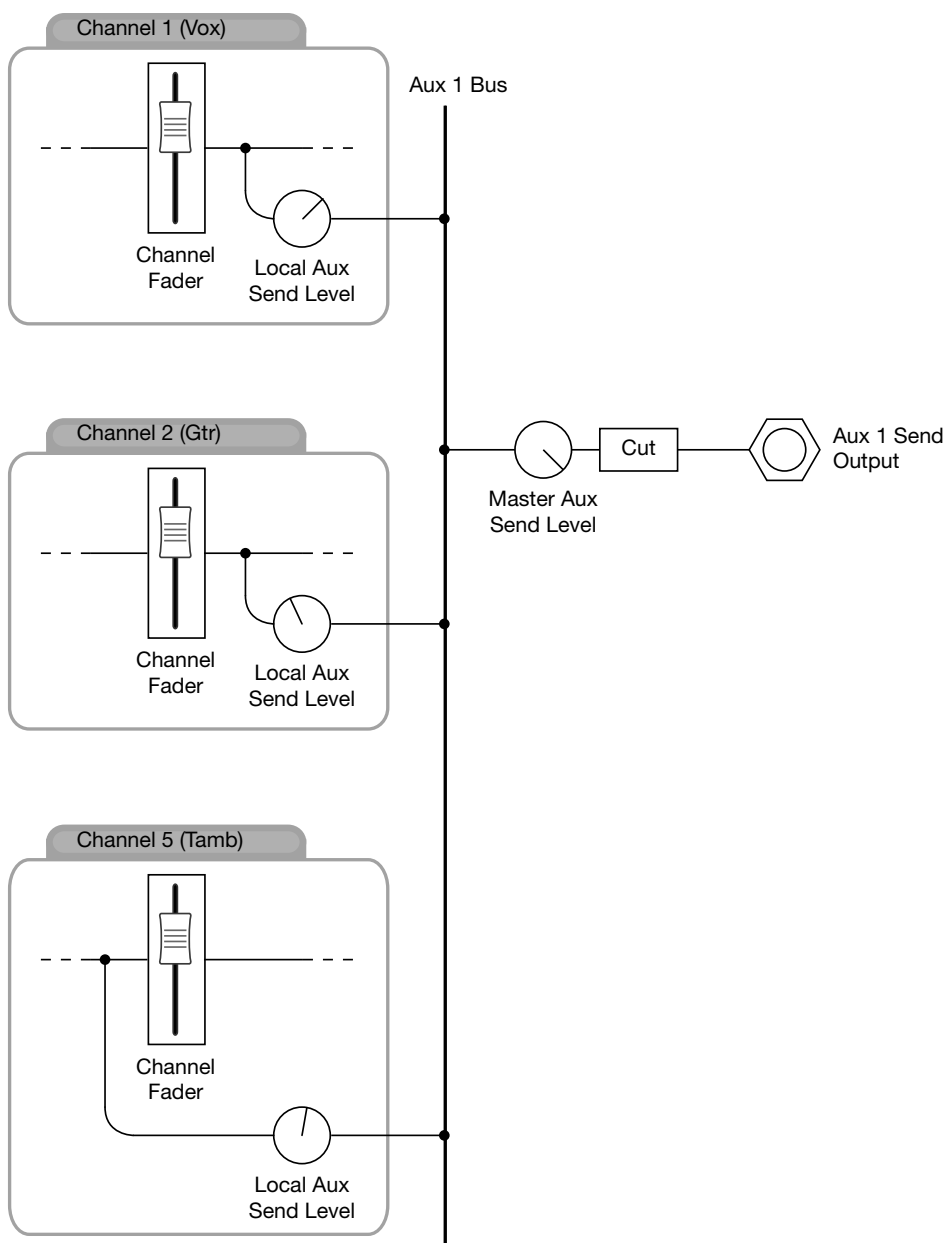
Figure 10.10 In this illustration, the signals from three channels are sent to the same auxiliary bus, each with its own send level settings. Note that for channels 1 and 2, a post-fader feed is selected, while for channel 5 there is a pre-fader feed.
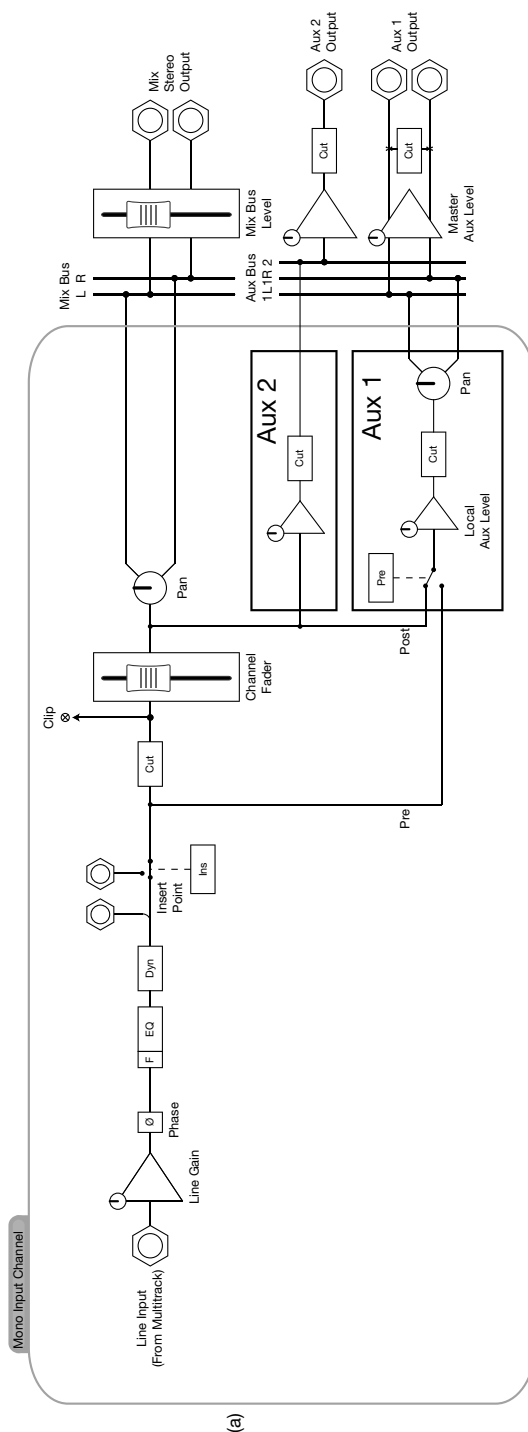
Figure 10.11 In step 5, auxiliary sends were added. The console has two auxiliary buses, one stereo and the other mono. As per channel, the local controls for aux send 1 are level, pre-/post-fader selection, pan, and cut. Aux 2 only has local level and cut controls. The master aux controls for both auxiliaries are level and cut.
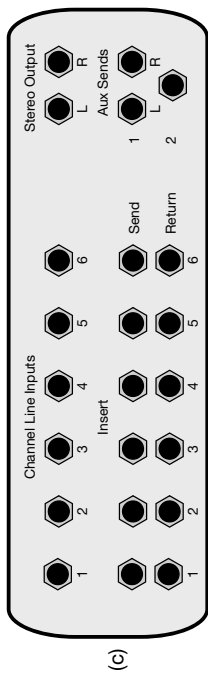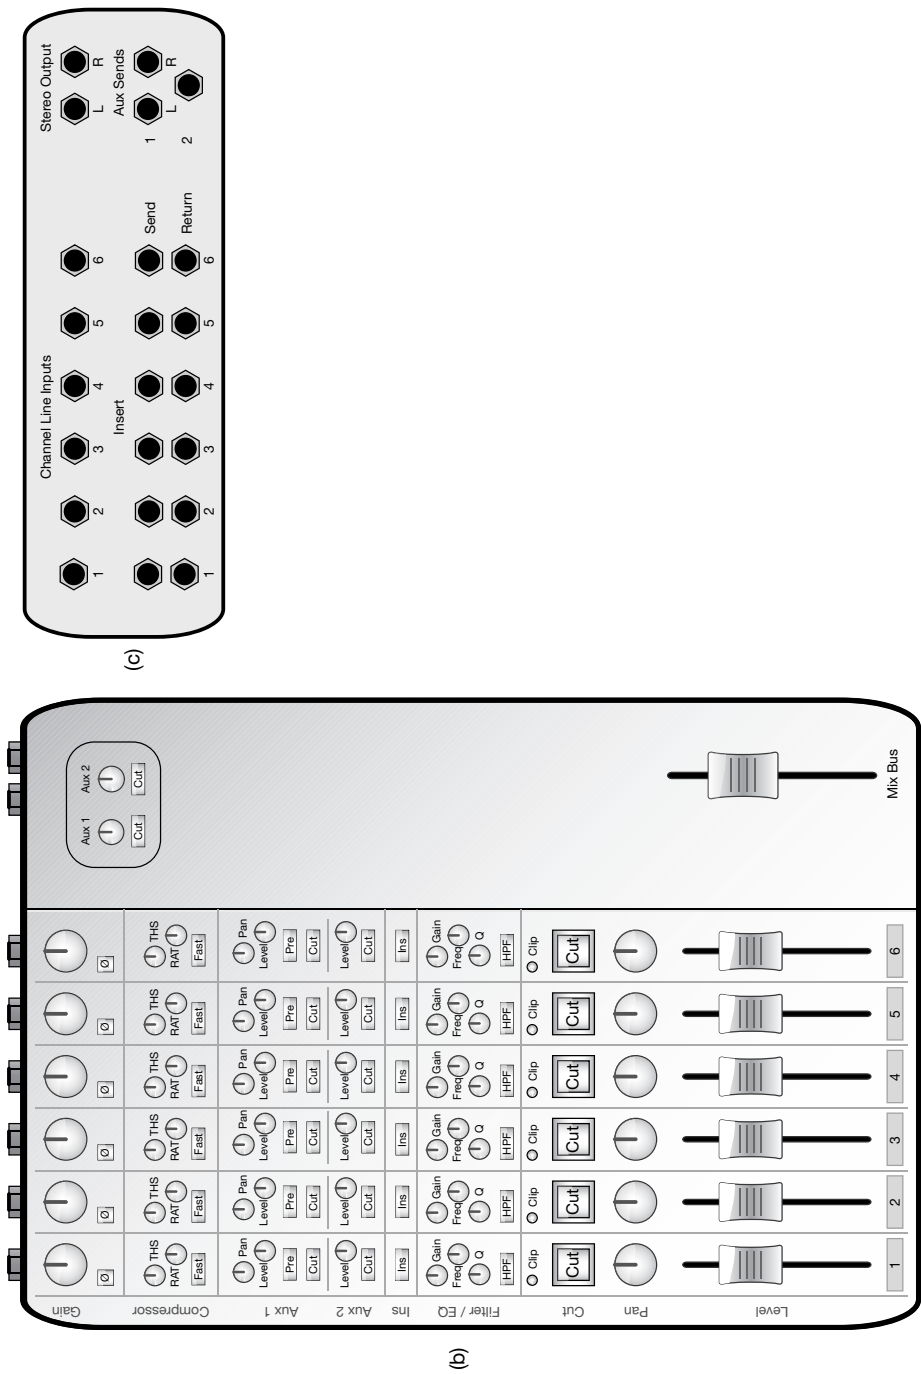
Figure 10.11 (*continued*)

Step 6: FX returns

So far, we have seen how we can send an audio signal to an effects unit. Let us consider now how the effect return can be brought back to the console. FX returns (sometimes called *aux returns*) are dedicated stereo inputs at the back of the desk that can be routed to the mix bus. All but the most basic desks provide some master controls to manipulate the returned effect signal. Level and mute are very common, while large-format desks might also have some basic tone controls. Figure 10.12 shows the addition of two stereo FX returns to our desk. In the layout in Figure 10.12b, a dedicated section provides level and mute controls for each return, and two pairs of additional sockets have been added to the rear panel in Figure 10.12c.

FX returns provide a very quick and easy way to blend an effect return into the mix, but even on large-format consoles they offer very limited functionality. For example, they rarely let us pan the returned signal, which might be needed in order to narrow down the width of a stereo reverb return. As an alternative, mixing engineers often return effects into two mono channels, which offer greater variety of processing (individual pan-pots, equalizers, compressors, etc.). Such a setup is shown in Figure 10.13.

When possible, effects are better returned into channels.

Groups

Rarely do raw tracks not involve some logical groups. The individual drum tracks are a prime example of a group that exists in most recorded productions; vocals, strings, and guitars are also potential groups. While a drum group might consist of 16 individual tracks, some groups might only consist of two—bass-mic and bass-direct, for example. During mixdown, we often want to change the level of, mute, solo, or process collectively a group of channels. The group facility allows us to do so. There are two types of groups: control and audio.

Control grouping

Control grouping is straightforward. We first allocate a set of channels to a group, so that moving one fader causes all the other faders in the group to move simultaneously. Cutting or soloing a channel also cuts or solos all other group channels. Some consoles define this behavior as *linking*, while grouping denotes a main-sub relationship—only changes to the main channel result in respective changes to the subs, but changes to subs do not affect other channels.

VCA grouping

VCA grouping is an alternative type of control grouping. In order for faders to respond to the movement of other faders, they must be motorized. Consoles without motorized faders can achieve something similar using a facility called VCA grouping. With VCA grouping, a set of master VCA group faders are located in the master section alongside the cut and solo buttons per group. Individual channel faders can then be assigned to a VCA group using a dedicated control per channel strip. Moving a master VCA group fader alters the level of all channels assigned to the VCA group (the channel faders will
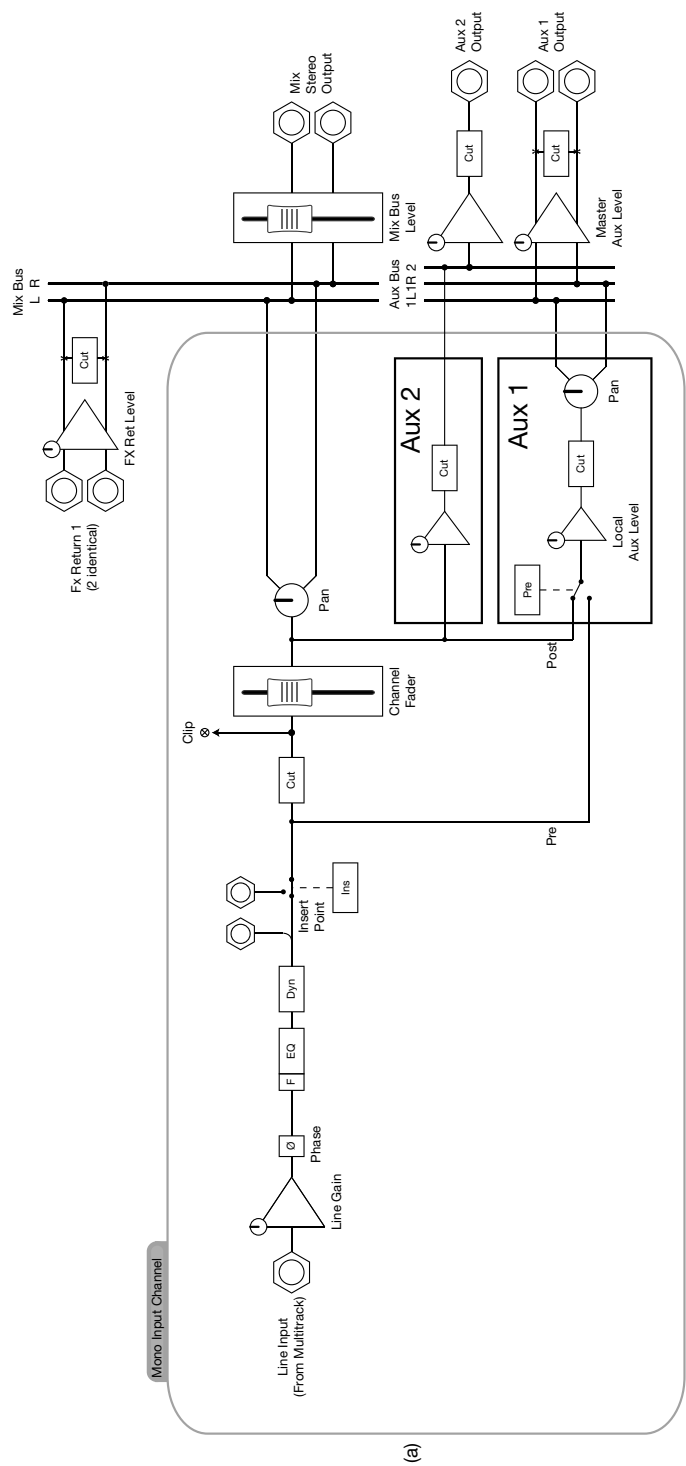
Figure 10.12 Step 6 involves the addition of two FX returns. There are two pairs of sockets on the rear panel; each FX return has level and cut controls that will affect the returned effect signal before it is summed to the mix bus.
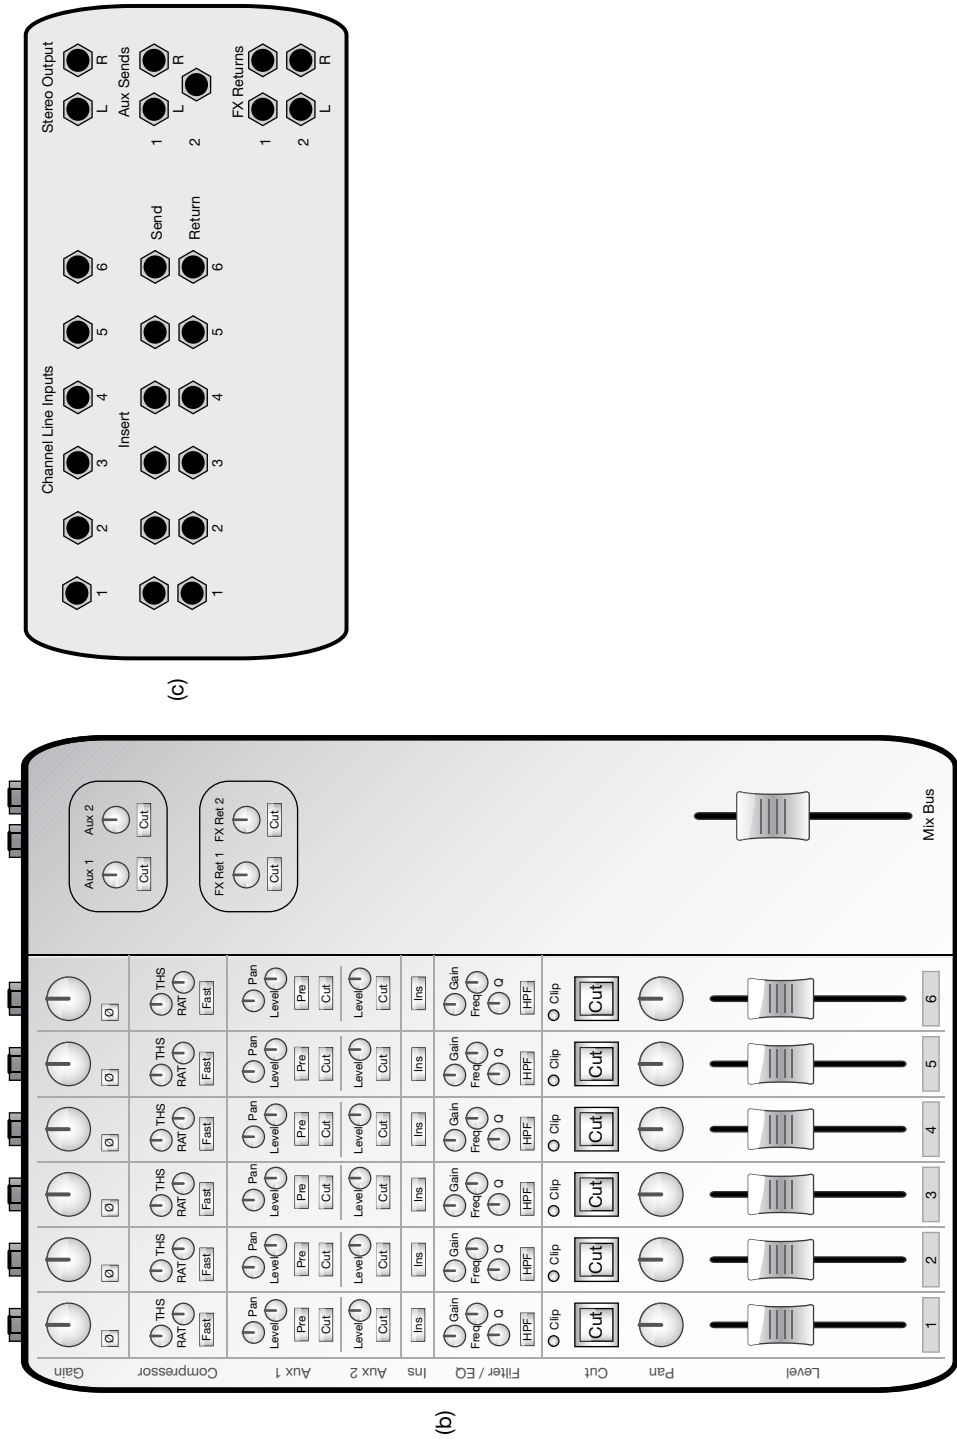
Figure 10.12 (*continued*)
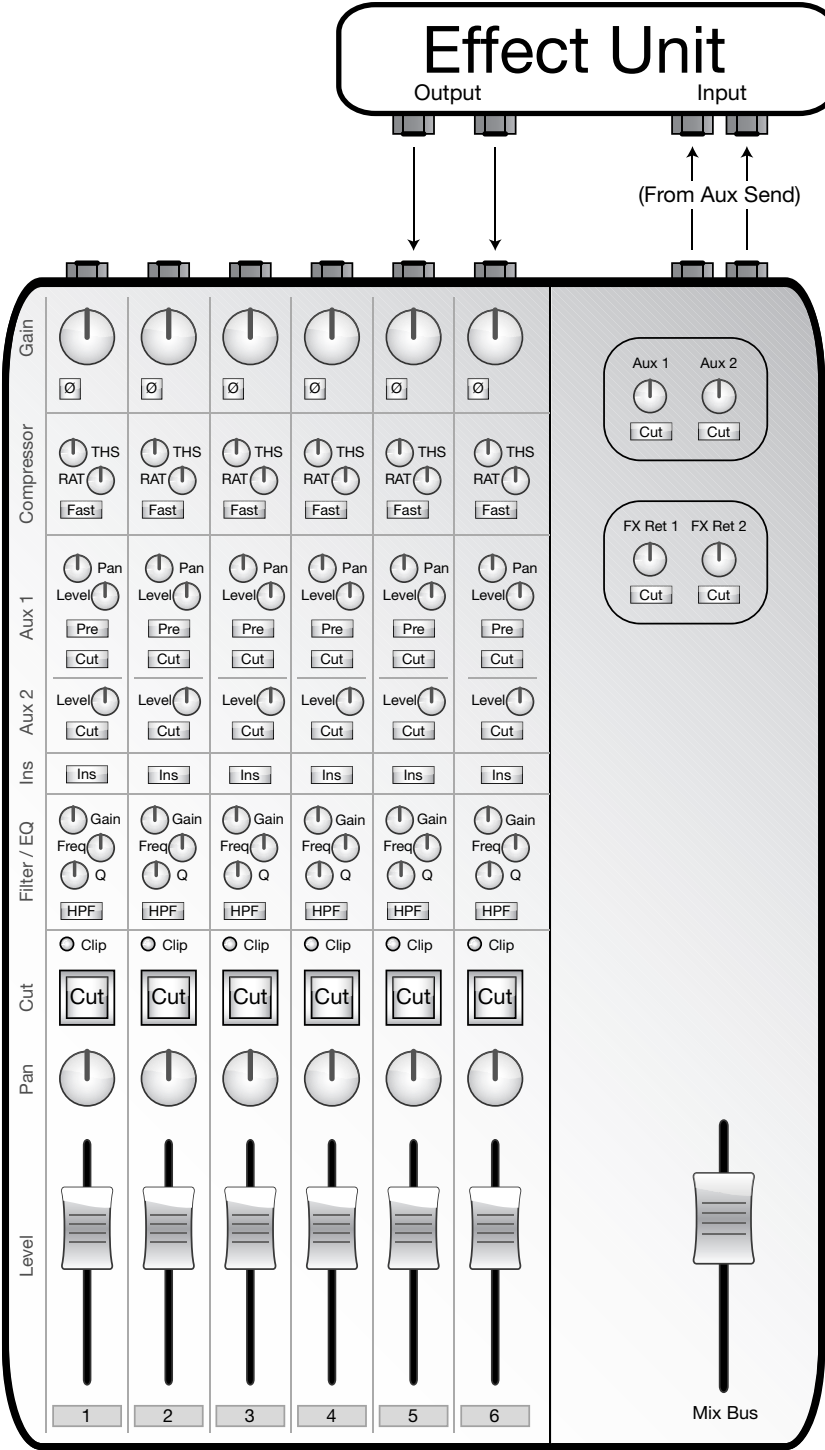
Figure 10.13 When there are free channels on the console, mixing engineers often bring back effect returns into channels rather than using FX returns. Channels offer more processing options compared to FX returns.

not physically move). Cutting or soloing a VCA group cuts or solos all the assigned channels. Figure 10.14 shows the layout of a small console with VCA groups.

The concept of VCA faders is explained in Chapter 13.
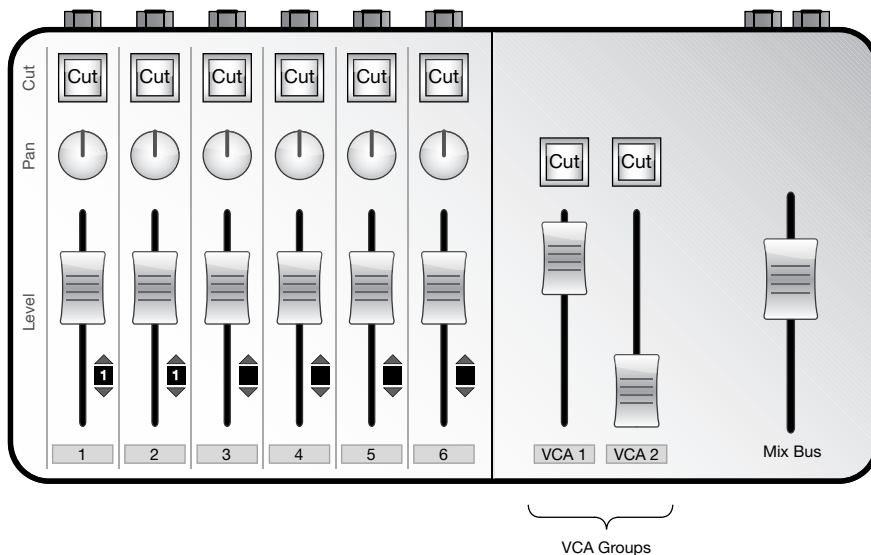
Figure 10.14 A basic console with two VCA groups. In the illustration above, channels 1 and 2 are assigned to VCA group 1. Moving the fader of VCA group 1 will alter the level of channels 1 and 2, although their faders will not move.

Control grouping is great when we want to *control* collectively a group of channels. In addition to the standard level, cut, and solo controls, digital consoles (and software mixers) sometimes enable the grouping of additional controls such as pan-pots, aux sends, or the phase-invert switches. However, control grouping does not alter the original signal path; therefore, it is impossible to *process* a group of channels collectively. This is often required; for example, in cases where we want to compress the overall drum mix. When collective processing is needed or if an analog console does not offer control grouping, audio grouping is the solution.

Audio grouping

To handle many signals collectively, a group of channels must first be summed to a group bus—a practice called *subgrouping*. The group signal (essentially a submix) can then be processed and routed to the mix bus. Different consoles provide different numbers of group buses, with multiples of eight being the most common. Each group bus is mono, but very often groups are used in stereo pairs made of consecutive odd and even groups—each represents left and right, respectively. For example, in groups 1–2, group 1 is left and 2 is right. The format *Channels:Groups:Mix-Buses* is commonly used to describe the number of these facilities on a console. For example, 16:8:2 denotes 16 channels, eight group buses, and two mix buses (i.e., one stereo mix bus).

(a)   (b)
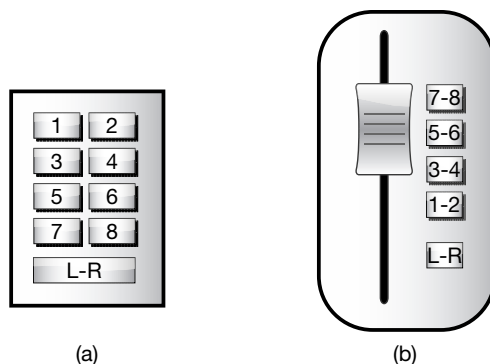
Figure 10.15 Routing matrices. (a) A typical arrangement for a routing matrix on a large console. Each group bus has a dedicated button. (b) A typical arrangement for a routing matrix on a smaller desk. Each button assigns the channel signal to a pair of groups, and the channel pan-pot can determine to which of the two the signal is sent.

Each channel can be assigned to different groups using a **routing matrix** that resides on each channel strip. A routing matrix is a collection of buttons that can be situated either vertically next to the fader or in a dedicated area on the channel strip (Figure 10.15), the latter being more common on larger desks, where there can be up to 48 groups. Routing matrices come in two formats—those that have an independent button for each group bus and those that have a button for each pair of groups. A pan-pot usually determines how the signal is panned when sent to a pair of groups. For example, if groups 1 and 2 are selected and the pan-pot is turned hard-left, the signal will only be sent to group 1.

As far as the signal flow goes, the signal sent to the group is a *copy* of the channel signal. We need a way to disconnect the original signal from the mix bus, since the copy sent to the group will be routed to the mix bus anyway (seldom do we want both the original and the copy together on the mix bus). To achieve this, each routing matrix has a mix assignment switch that determines whether or not the original channel signal feeds the mix bus. The copy is taken *post the channel pan-pot* (and thus post the channel fader), and therefore any processing applied to the channel signal (including fader movements) affects the copy sent to the group. It should be made clear that with audio grouping, we have full processing control over each individual channel and on the group itself. We can, for example, apply specific compression to the kick, different compression to the snare, and different compression still to the drum mix.

Group facilities fall into two categories: master groups and in-line groups. Consoles that offer **master groups** provide a dedicated strip of controls for each group bus on the master section. Each strip contains a fader that sets the overall level of the group, and a few additional controls such as a solo button or a mix-bus assignment switch. The latter is required, as during recording or bouncing the group buses are not normally assigned to the mix bus. The mix-bus assignment switch also acts as a mute control since it can disconnect the group signal from the mix bus. To enable the processing of the group signal, each group offers an insertion point.

Figure 10.16a shows the signal flow of a console with master audio groups. Note that the signal to the group bus is a copy taken post-pan. Each group has an insert point (in
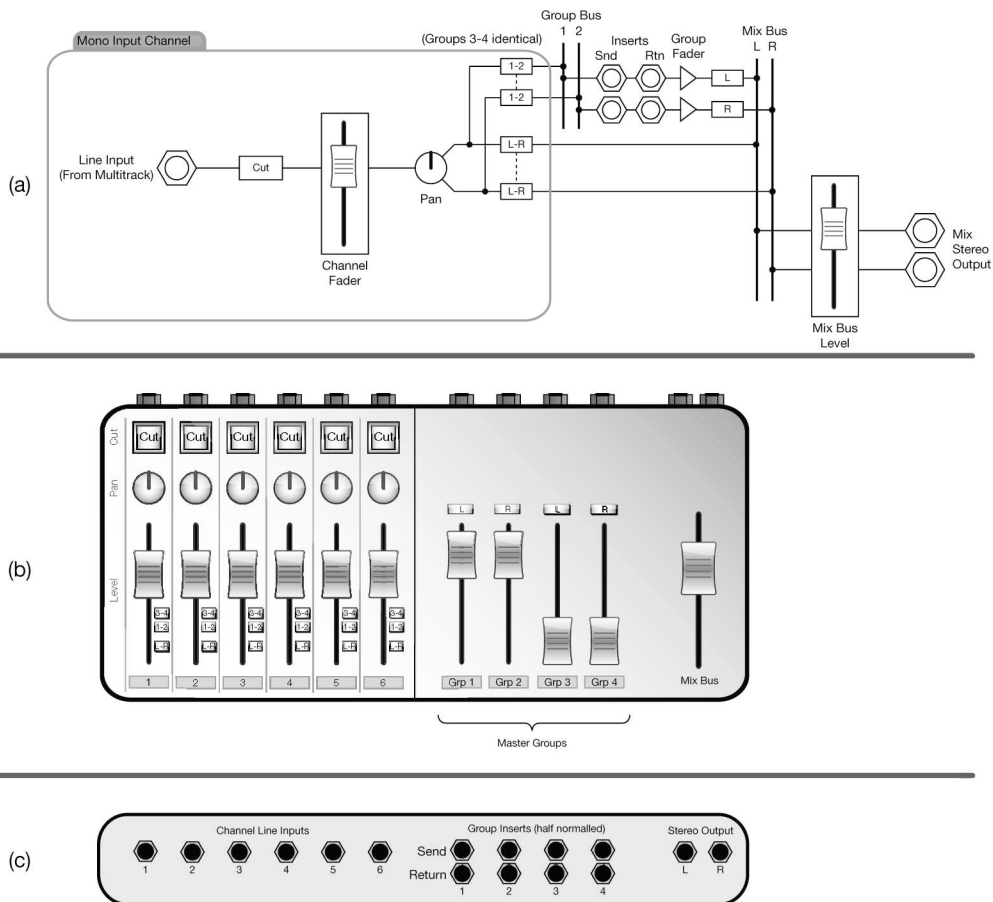
Figure 10.16 A basic console with four master audio groups. In this illustration, channels 1–2 are assigned to groups 1–2; the rest of the channels are assigned to the mix bus and so are groups 1–2.

this case half-normaled, like on a standard patchbay), a fader, and a mix-bus assignment switch. With this console, odd groups can only be assigned to the left, and even groups to the right—effectively forcing the groups to work as stereo pairs. Most consoles also have a button to assign each group to both channels of the mix bus. In Figure 10.16b, we can see the routing matrix next to each channel fader, which involves an assignment button to the mix bus and to each pair of groups. There are also four master group strips; each contains a fader and an assignment switch to either the left or right mix buses. Figure 10.16c shows the addition of the group insert sockets to the rear of the console.

In-line groups are only found on large-format in-line desks. These desks still have a routing matrix per channel, but they do not offer any master group strips. Instead, each channel strip hosts its associated group (for instance, channel 1 hosts group 1, channel 2 hosts group 2, and so forth). Once hosted, the group signal is affected by all the channel components, including the fader, pan-pot, onboard equalizer, compressor, and insert point. These provide much more functionality than any master group might. But this functionality comes at a price—each channel can only be used to either handle a
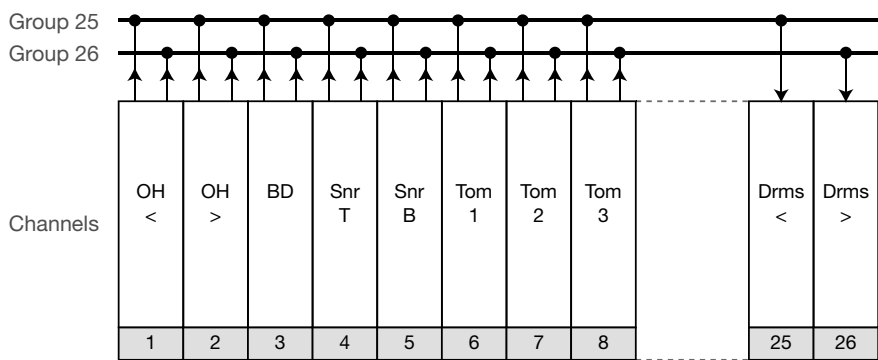
Figure 10.17 Basic schematic of audio groups. Channels 1–8 are routed to groups 25–26, which are then fed into their associated channels.

track from the multitrack or host its associated group; it cannot do both. In order to use in-line groups, we must have more channels than tracks. On most large consoles, this is not an issue—if the multitrack has 24 tracks, they will feed channels 1–24, while channels 25–48 will be available as in-line groups. A possible scenario would involve disconnecting all the drum channels (say, 1–8) from the mix bus and routing them to groups 25–26. Channels 25–26 would then be used to host groups 25–26. To route the two groups to the mix bus, all we would need to do is select the mix bus on the routing matrices of the hosting channels (Figure 10.17).

A group button on the input section of each channel determines what signal feeds the channel path: either the multitrack return signal (the track coming back from the multitrack via the line input) or the associated group signal (which turns the channel into its associated group host). In the case of the drum scenario illustrated above, it will be pressed on channels 25 and 26. Pressing the group button on a channel that has its associated group selected on the routing matrix will cause feedback (the channel will be routed to the same group that feeds it).

Figure 10.18 illustrates a hypothetical desk, as it involves in-line groups on a split desk. Yet it demonstrates well the concept of audio groups. The signal flow in Figure 10.18 shows that pressing the group button will feed the associated group signal into the channel path before the phase button. All the groups in this case are just being sent back to their associated channels. Figure 10.18b reveals a dedicated routing matrix per channel with an assignment button per group (and one button for the mix bus). Also, the group button was added next to the phase button.

Bouncing

The concept of bouncing is explained in Appendix A. So far, we have only discussed groups in relation to mixing, but groups are an essential part of the recording process as they provide the facility to route incoming microphone signals to the multitrack recorder. Every console that offers group buses also has an output socket per group bus. The group outputs are either hardwired or can be patched to the track inputs of the multitrack recorder. The process of bouncing is similar to the process of recording, except that instead of recording microphone signals, we record a submix.
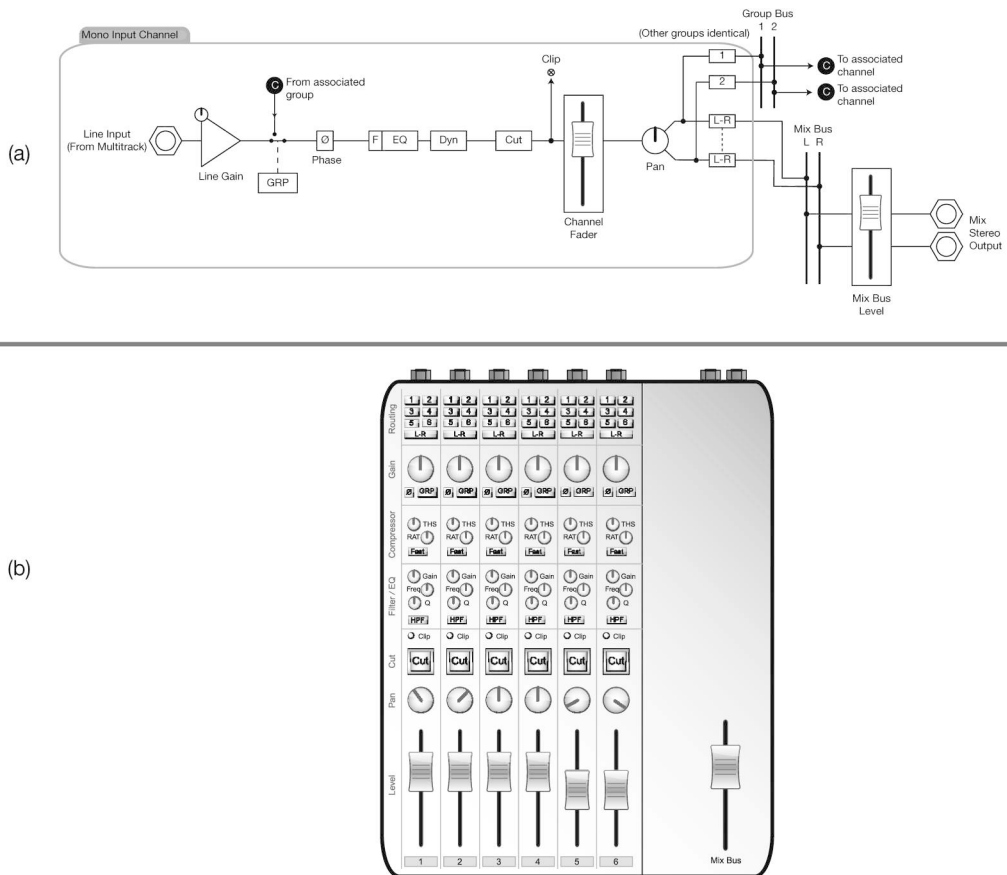
Figure 10.18 A console with in-line group facilities. Channels 1–3 are routed to groups 5–6. The pressed group buttons on channels 5–6 bring the group signal into these channels, which are the only ones assigned to the mix bus.

In order to bounce something, all we have to do is route the channels of our submix to a group and make sure that the group output is sent to an available track on the multitrack recorder. Since most submixes are stereo, this process usually involves a pair of groups.

Figure 10.19 shows the console in Figure 10.18 but with the additional group outputs. If you remember, the less-than-optimum level is one of the most common issues with bouncing, and setting the group level is one way to bring the bounced signal to its optimal level. Figure 10.19a shows that the group signal ends up in a group output socket, but only after it has passed a level control. One question that might arise is: Where exactly on the desk do we put these group-level controls? We could place them on the master section, but since consoles with in-line groups have as many groups as channels (commonly up to 48 groups) and since these channels and groups are associated, it makes sense to put the group level on each channel strip. As can be seen in Figure 10.19b, the group-level control has been added just below the group button. Figure 10.19c illustrates the addition of the group output sockets, as expected, aligned with the channel inputs.
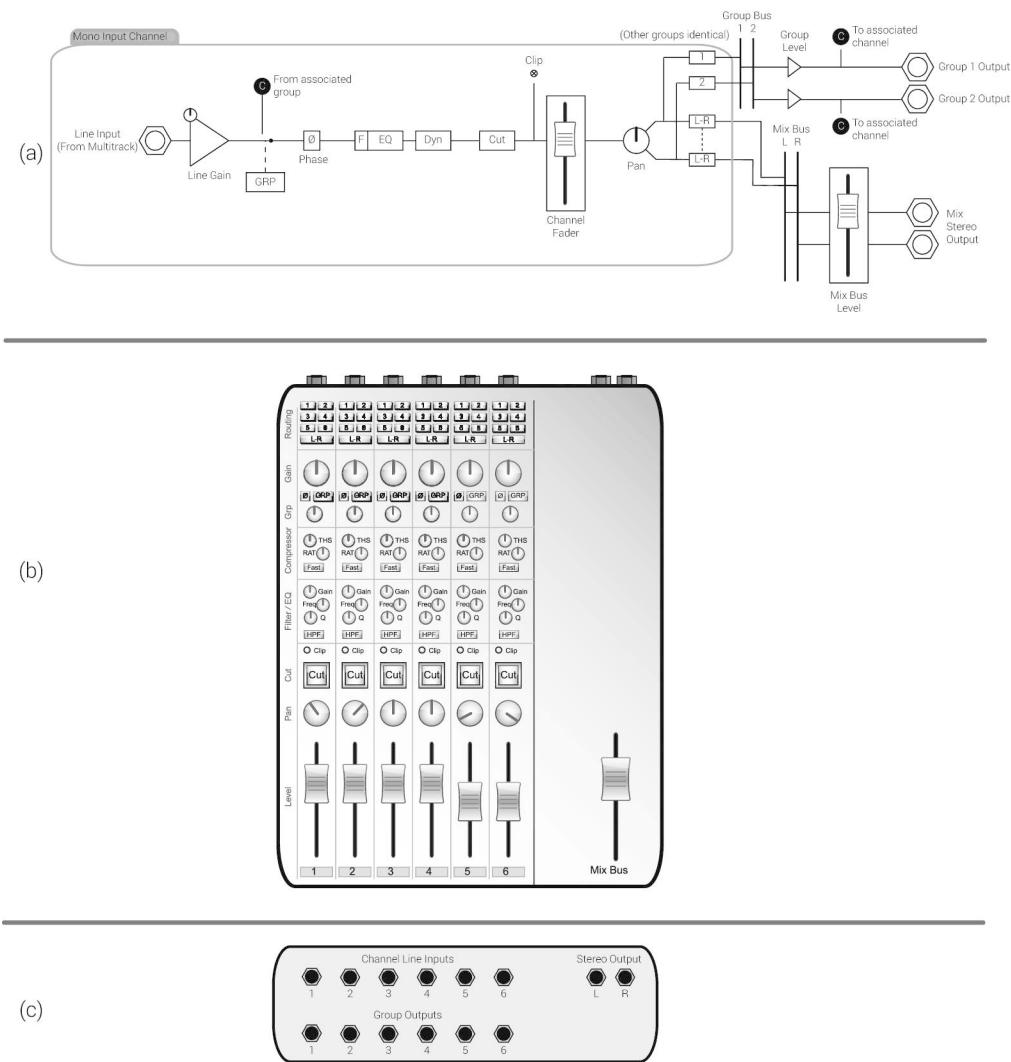
Figure 10.19 A console with physical group outputs, and group-level controls.

In-line consoles

The recording aspects

It is impossible to discuss the merits of in-line consoles without talking about recording, but discussion of the latter will be kept brief. During a recording session, the desk accommodates two types of signal: (a) the live performance signals, which are sent via the groups to be recorded on the multitrack; and (b) the multitrack return signals, which include any previously recorded instruments. On a split desk, which we have discussed so far, some channels handle the live signals while others handle the multitrack return signals. But a single channel cannot handle two signals at the same time. Figure 10.20 shows the layout of a split desk during recording.
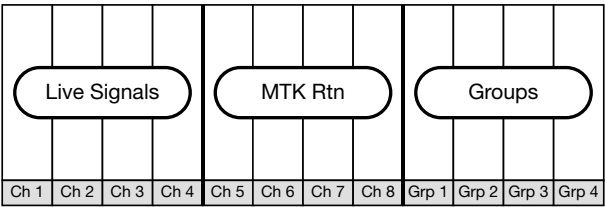
Figure 10.20 An eight-channel, four-group desk as might be utilized in a recording situation. Channels 1–4 will handle live signals coming through the microphone or line inputs, channels 5–8 will handle the signals returning from the multitrack via the line inputs, and the master group will be used to route the incoming signals to the multitrack.

Split designs can be large. For example, a 24-track recording would require a 48-channel console—24 channels for live signals and 24 for multitrack returns. If the console also includes 24 master groups, we end up with 72 strips. Size aside, each channel strip costs a good amount of money. However, we rarely require all the expensive features on all channels—we either process the live signals or the multitrack return signals, seldom both.

Going in-line

The way in-line designs solve the size issue is by squeezing 48 channels into 24 I/O modules (referred to as channel strips on a split desk). This is done by introducing additional signal path to the existing channel path on each I/O module. While the channel path is used for live signals, the new path is used for multitrack return signals. Since it is the latter that we monitor (both in the control room and in the live room), the new path is called the *monitor path*. The additional requirements for each I/O module include a line input socket for the multitrack return signal and a fader (or pot on smaller desks) to control the level of the monitor path signal. Large in-line consoles are easily identified by the two faders that reside on each I/O module (Figure 10.21). There will still be only one equalizer,

Figure 10.21 The SSL Origin at Abbey Road Institute, London. This photo clearly shows the two faders per I/O module. Depending on the desk configuration, one fader controls the level of the monitor path, the other the channel path.

compressor, aux send, and insert point per I/O module, but each facility can be switched to either the channel or monitor path (console permitting).

So we can condense a 48-channel split desk into a 24-channel in-line desk, but what about the additional 24 master groups? If we exchange them for 24 in-line groups (as in Figure 10.18), we can integrate those groups into our 24 I/O modules. We will then end up with 24 I/O modules, each with a channel path, monitor path, and a group, all in-line (Figure 10.22).
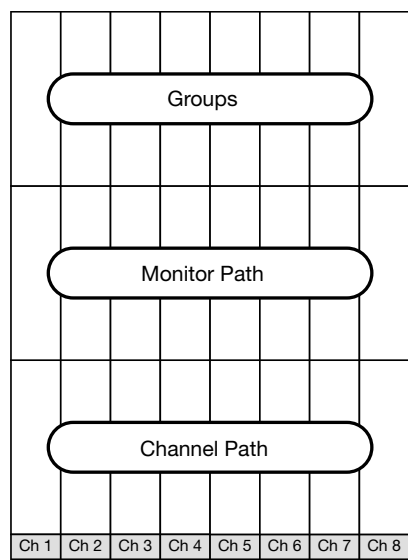
Figure 10.22 An eight-channel in-line console.

A 24-track studio might benefit from the in-line concept, but if the studio has a small live room—where only eight microphones may be used simultaneously—it would only need eight groups on its console. A specific design combines a channel and monitor path per I/O module, with a master instead of in-line groups. Figure 10.23 illustrates this concept, while Figure 10.24 shows such a console.
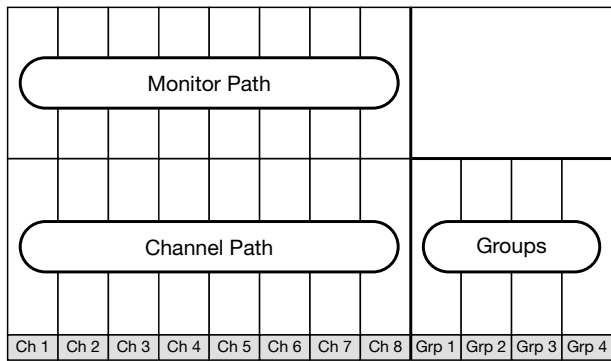
Figure 10.23 An eight-channel, four-group console.
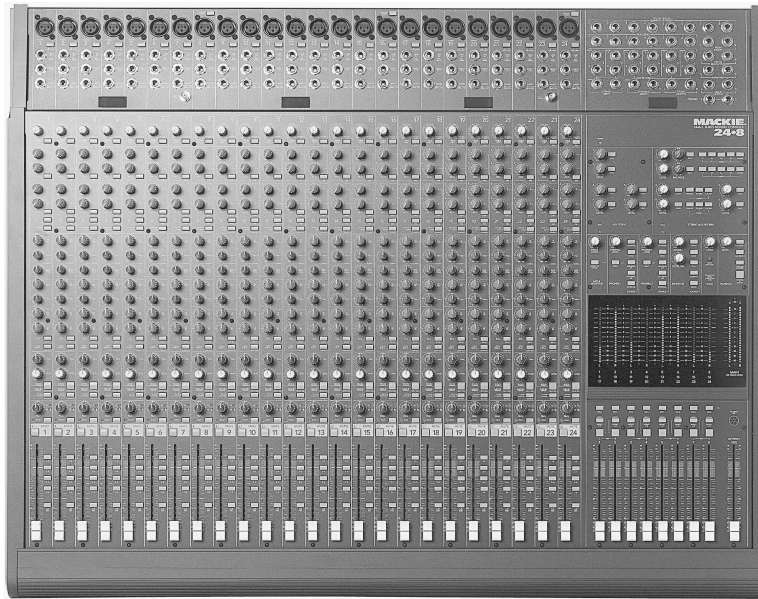
Figure 10.24 The Mackie Analog 8-Bus. Each I/O module has separate level and pan controls per path. In the case of the 8-Bus, the level of one path is controlled using a fader, while the level of the other is controlled using a pot (the white row left of the meters). The eight master group faders can be seen on the right (between the channel faders and master fader).

In-line consoles and mixing

The in-line design makes consoles smaller and cheaper, but as effective for recording. When it comes to mixing, they make things slightly more complex than the split design. The first thing to understand is that a 24-channel console is able to accommodate 48 input signals as each I/O module has both channel and monitor paths. In a 24-track studio, this means the addition of 24 free signal paths during mixdown. As already mentioned, the different processing and routing facilities can only be switched to one path at a time. The catch is that not all facilities can be switched to the monitor path—the insert point, for example, might be fixed on the channel path. Generally, the channel path is the "stronger" path—providing full functionality—while the monitor path is more limited. It is therefore reasonable to use the channel path for the main mix of multitrack return signals and utilize the monitor path for a variety of purposes. Under normal circumstances, the monitor path is used for the following:

• **Effect returns**—effect returns can be easily patched to the monitor path via its line input. We can, for example, send a guitar from the channel path to a delay unit and bring the delay back to the monitor path on the same I/O module.
• **Additional auxiliary sends**—even the most respectable large-format consoles can be limited to five auxiliaries, only one of which might be stereo. While this is sufficient for recording, it can restrict effect usage during mixdown—what if we need more than five effects? By having a copy of the channel signal in the monitor

path, we can route it to a group, which is then routed to an effects unit. In this scenario, the monitor path level control acts as a local aux send, while the group level acts as master auxiliary level.
• **Signal copies**—in various mixing scenarios, a signal duplicate is needed and the monitor path is the native host for these duplicates. One example is the parallel compression technique.

The monitor section

In addition to the global sections we have covered so far, the master section may also contain subsections dealing with global configuration, monitoring, metering, solos, cue mix, and talkback. Out of these, the monitor section is the most relevant for mixdown. Needless to say, the selection of features is dependent on the actual make. The following section is limited to the most common and useful features.

The monitor output

The imaginary console we have built so far has had a stereo mix output all along. The mix output, conventionally, only outputs the mix-bus signal and should be connected to a two-track recorder. Although most of the time we monitor the mix bus, sometimes we want to listen to an external input from a CD player, a soloed signal, an aux bus, or some other signals. To enable this, consoles offer *monitor output* sockets, which, as you might guess, are connected to the monitors. There is also a separate gain control, usually a pot, to control the overall monitoring level.

Figure 10.25 shows the addition of the monitor output facilities to the basic console in step 1. In Figure 10.25a, both the mix bus and an external input (2TK input or two-track input) are connected to the monitor output. A control circuit cuts all the inactive monitor alternatives. The monitor gain is determined by the *CR Level* pot (control room level). It is worth noting that the mix level will affect the monitor level, but not the other way around. This lets us alter the monitor level without changing the signal level sent to a two-track recorder via the mix output. If we want to fade our mix, we must use the mix level. Figure 10.25b illustrates the addition of the monitor gain control, labeled "monitor level." At the rear of the desk in Figure 10.25c, a pair of monitor output sockets have been added, as well as a pair of external input sockets.

Additional controls

Figure 10.26 illustrates a master monitor section, very similar to that found on a large-format console. The various controls are:

• **Cut**—cuts the monitor output. Used in emergencies, such as sudden feedback or noise bursts. Also used to protect the monitors from predictable clicks and thumps caused by computers being turned on and off, or patching.
• **Dim**—attenuates the monitor level by a fixed number of dBs, which is sometimes determined by a separate pot (around 18 dB is common). The monitor output will still be audible, but at a level that enables easier communication between people in the room or on the other end of a telephone.
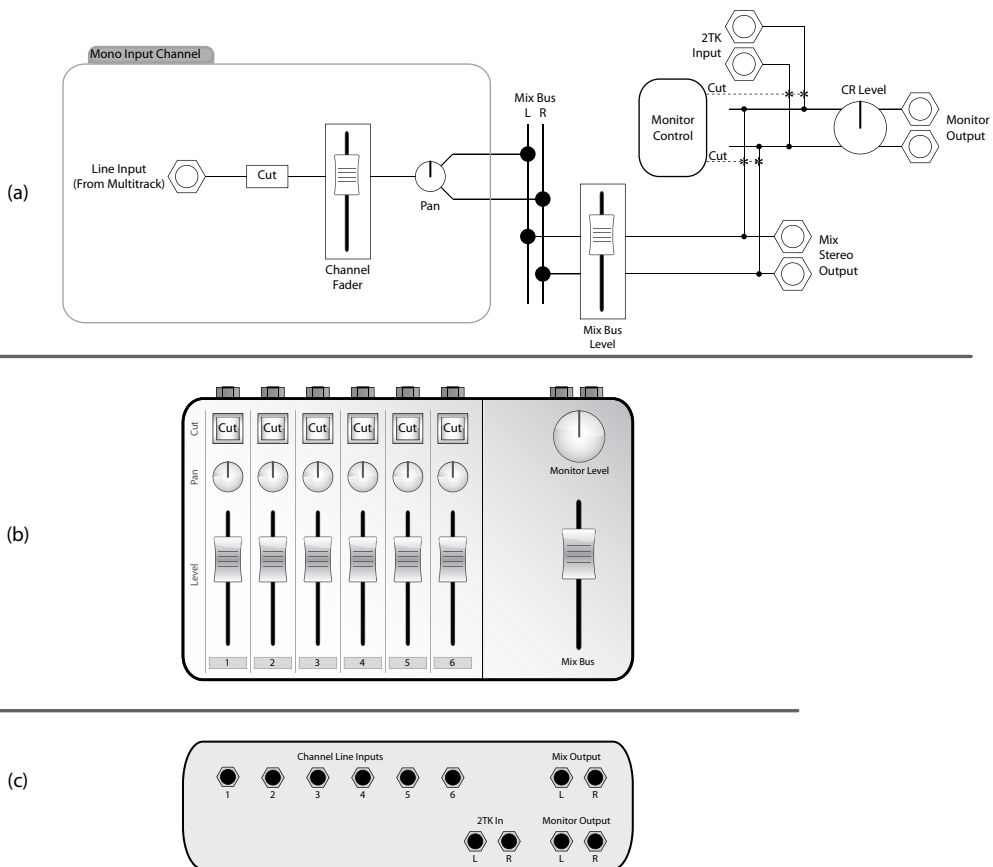
Figure 10.25 A basic console with monitor output facilities.

- **Mono**—sums the stereo output into mono. Enables mono-coherence checks and can be used as an aid when dealing with masking.
- **Speaker selection**—a set of buttons for switching between the different sets of loudspeakers; that is, the full-range, near-fields, and mini-speakers. It is also possible to have an independent switch that toggles the subwoofer on and off. A console with such a facility will have individual outputs for each set of speakers.
- **Cut left, cut right**—two switches: each cuts its respective output channel. Used during mix analysis and to identify suspected stereo problems; for example, a reverb missing from the left channel due to a faulty patch lead.
- **Swap left/right**—swaps the left and right monitor channels. Swapping the mix channels can be disturbing since we are so used to the original panning scheme. Also, in a room with acoustic stereo issues (due to asymmetry, for example), pressing this button can result in altered tonal and level balances, which are even more disturbing. This button is used to check stereo imbalance in the mix. If there is such an imbalance, pressing this button might make us feel like turning our head toward one of the speakers; if it is unclear toward which speaker the head should turn, the stereo image is likely to be balanced. Image shifting can also be the result of problems in the monitoring system, such as when the vocals are panned center but appear to come from an off-center position. If by pressing this button the vocals
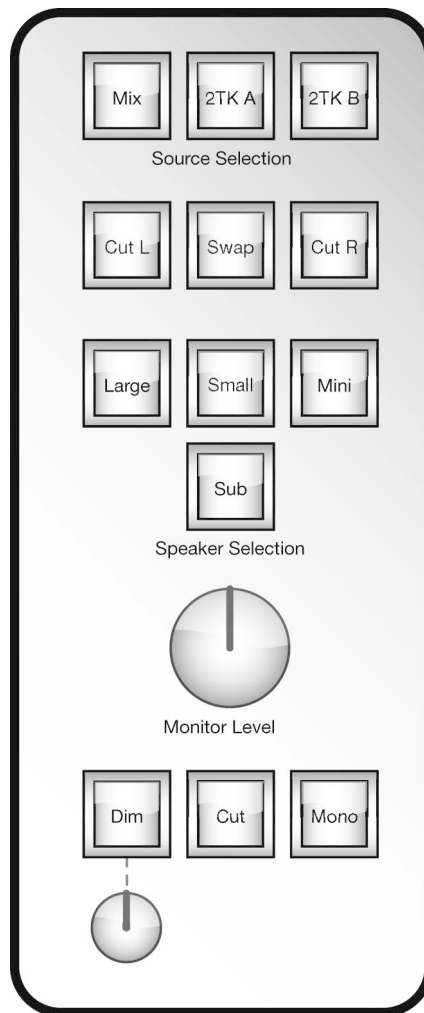
Figure 10.26 A master monitor section.

remain at the same off-center position, one speaker might be louder than the other or one of the speaker cables (or amplifier channels) might be attenuating the channel signal. If the vocal position is mirrored around the center, then it is clear that the shifting is part of the mix itself.

• **Source selection**—determines what feeds the monitor outputs. Sources might include the mix bus, external inputs, or an aux bus. Sources might be divided into internal and external sources, and additional controls let us toggle between the two types.

Solos

Solo modes determine how the solo facility operates; that is, what exactly happens when a channel is soloed. There are two principal types of solos: destructive and nondestructive. Nondestructive solos can be PFL, AFL, or APL (the meanings of which will be covered shortly). The hierarchy of the different solo modes is as follows:

- destructive in-place; and
- nondestructive:
  - PFL; and
  - AFL or APL.

Large consoles often support more than one type of solo mode; for example, a desk might support destructive solo, PFL, and APL. A console can either support AFL or APL solo, but rarely both. Often manufacturers use the term *AFL* for a facility that is essentially APL.

The active solo mode can be selected through a set of switches that reside on the master section. In certain circumstances (discussed soon), a desk will toggle momentarily to nondestructive mode.

Destructive solo

Destructive solo might also be called *in-place*, *destructive in-place*, *mixdown solo*, or *SIP*, which stands for *solo in-place*. In destructive solo mode (Figure 10.27), whenever a channel is soloed, all the other channels are cut. Therefore, only the soloed channel (or channels) is sent to the mix bus (which is still monitored). It should be clear that both channel level and panning information are maintained with this type of solo.
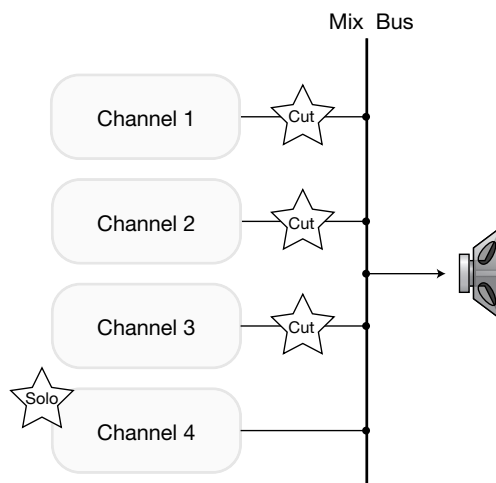
Figure 10.27 Destructive in-place solo. (It is worth knowing that in practice the individual channels are cut by internal engagement of their cut switch.)

Nondestructive solo

In nondestructive mode (Figure 10.28), no channel is being cut. Instead, the soloed channels are routed to a solo bus. As long as solo is engaged, the console automatically routes the solo bus to the monitors, cutting momentarily the existing monitor source (normally the mix bus). With nondestructive solo, the mix output remains intact as all the channels are still summed to the mix bus. PFL (pre-fade listen), AFL (after-fade listen),
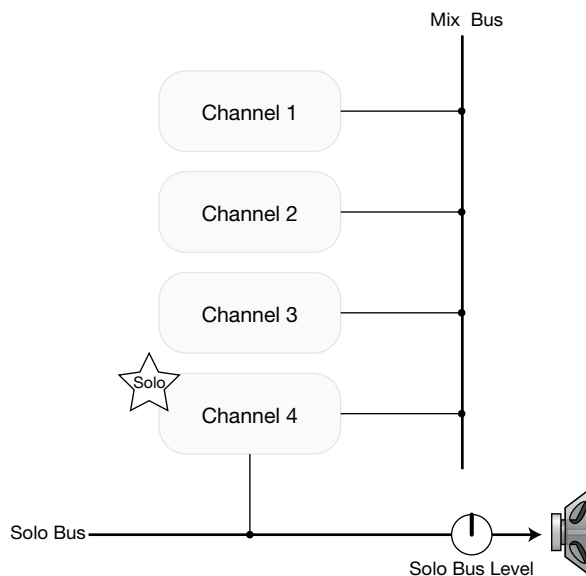
Figure 10.28 Nondestructive solo.
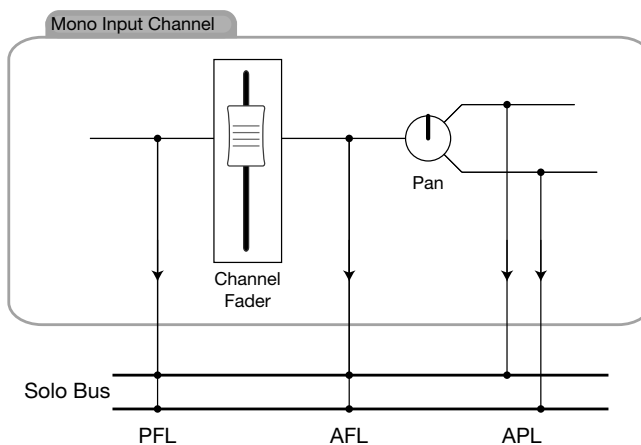
Figure 10.29 Depending on the console and the global solo mode, soloed signals can be sourced from one of three different points along the channel path.

and APL (after-pan listen) signify the point along the channel signal path from which the signal is taken (Figure 10.29). PFL takes a copy before the channel fader and pan-pot; therefore, both mix levels and panning are ignored. AFL takes a copy after the fader but before pan, so it maintains mix levels but ignores panning. APL takes a copy after both fader and pan, so both level and panning are maintained. Only a desk that provides APL solo requires a stereo solo bus; for PFL and AFL solos, a mono bus will suffice.

A solo-bus level control is provided on the master section. It is recommended to set this control so that signals are neither boosted nor attenuated once soloed. Essentially, this involves leveling the solo bus with the mix-bus level. Any other setting and the mixing engineer might be tempted to approach the monitor level following soloing;

then, when solo is disengaged, the monitored mix level will either drop or rise, which might again lead to alteration of the monitor level. The greatest danger of this is constantly rising monitor levels.

Solo safe

Solo safe (or *solo isolate*) provides a remedy to problems that arise as part of the destructive solo mechanism. More specifically, it prevents channels flagged as solo safe from being cut when another channel is soloed. A good example of channels that should be solo safe would be those hosting an audio group. If soloing a kick subsequently cuts the drum group to which it is routed, the monitors output null (Figure 10.30a). To prevent them from being cut, the channels hosting the drum group are flagged as solo safe. The same null can occur if the audio group channels are soloed since this will cut the source channels, including the kick (Figure 10.30b). Unfortunately, a console cannot determine the individual source channels of an audio group, so in order to prevent them from being cut the console would automatically engage into nondestructive solo mode when solo-safe channels are soloed. Along with audio groups, effect returns are also commonly flagged as solo safe.
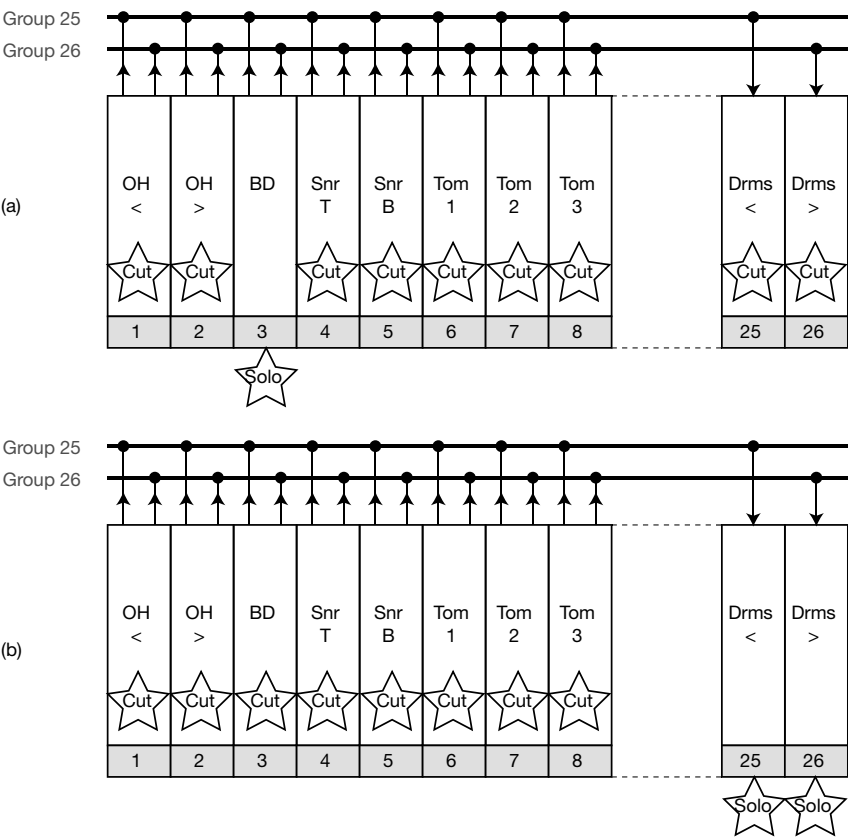
Figure 10.30 Two problematic scenarios with destructive solos and audio grouping. (a) Soloing a source channel will cut the group it is routed to. (b) Soloing the group hosting channels will cut the source channels.

Which solo?

Destructive solo is considered unsafe during recording. Recording during mixdown only happens in two situations: when we bounce and when we commit our mix to a two-track recorder. In these cases, any destructive solo will be imprinted on the recording, which is seldom our intention. With the exception of this risk, destructive solo offers many advantages over nondestructive solo. First, if the console involves a mono solo bus, all soloed signals are monitored in mono. This is highly undesirable, especially for stereo tracks and effect returns (e.g., overheads or stereo reverb). With destructive solo, we listen to the stereo mix bus, so panning information is maintained. Another issue with nondestructive solo has to do with unwanted signals on effect returns. Consider a mix where both the snare and the organ are sent to the same reverb unit, and in order to work on the snare and its reverb we solo both. In this scenario, the reverb return still carries the organ reverb, since nothing prevents the organ from feeding its aux send. By soloing in destructive mode, the organ channel will be cut and so will its aux feed to the reverb. Destructive solo also ensures that the soloed signal level remains exactly as it was before soloing. In nondestructive solo, soloed signals might drop or rise in level, depending on the level of the solo bus.

Destructive solo is favored for mixdown.

The only time PFL solo comes in handy during mixdown is when we want to audition the multitrack material with faders down.

Correct gain structure

Correct gain structure is a tactic that helps us to reduce unwanted noise, thus improving the overall signal-to-noise ratio of our mix. Analog components are noisy—microphones, preamps, compressors, equalizers, pan-pots, and analog tapes are just a few examples of components that add their inherent noise to the signal. It would be fair to say that from the point where an acoustic performance is converted into an electrical signal to the point where the same signal flows on the mix bus as part of the final mix, hundreds of analog components have added their noise. High-quality gear (such as large-format consoles) is built to low-noise specifications, but some vintage gear and cheap equipment can add an alarming amount of noise. Whatever quality of equipment is involved, simple rules can help us to minimize added noise.

The principle of correct gain structure is simple: never boost noise. Say, for example, that an equalizer has inherent noise at –70 dB, and we feed into it a signal with a peak at –20 dB. The resultant output SNR (signal-to-noise ratio) will be 50 dB. If we then boost the equalizer output by 20 dB, we boost both the noise and the signal, ending up with the signal level at 0 dB but noise at –50 dB. The SNR is still 50 dB (Figure 10.31a). Now let us consider what happens if we boost the signal *before* the equalizer. The peak signal of –20 dB will become 0 dB as before, but the noise added by the equalizer would still be –70 dB. The output SNR will be 70 dB (Figure 10.31b).

So when boosting the level of something, we also boost the noise of any preceding components. We can do a lot worse than attenuating something and then boosting it again. Figure 10.32 illustrates a signal at 0 dB going through two gain stages; one
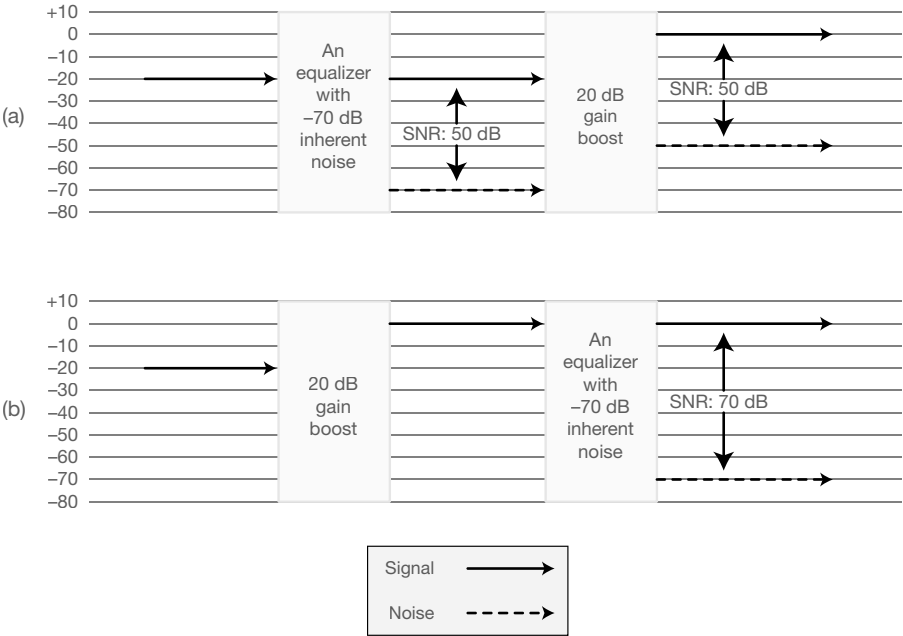
Figure 10.31 A demonstration of gain application and its effect on SNR. (a) The gain is applied after the noise is added, resulting in an output signal with 50 dB SNR and noise at –50 dB. (b) The gain is applied before the noise is added, resulting in higher SNR and lower noise floor. The noise level of both the original signal and the gain stage are omitted from this illustration to keep it simple. We can assume that these are too low to be displayed on the given scale.
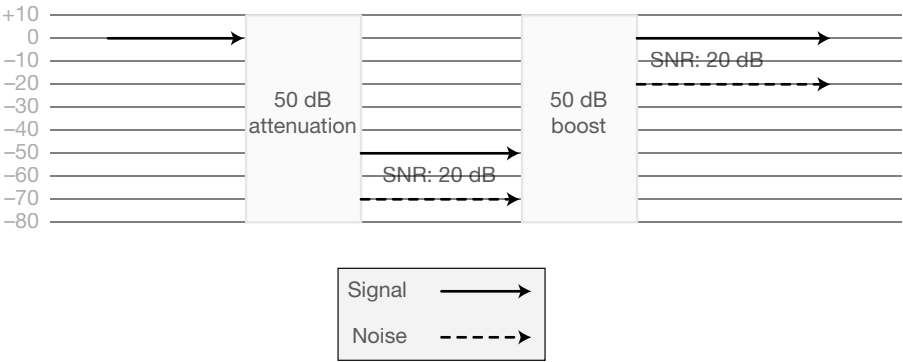

Figure 10.32 The penalty in SNR when boosting after attenuating. An incoming signal will leave the system at the same level, but the noise from the first gain stage will be boosted by the second gain stage, resulting in an output noise level of –20 dB and an SNR of 20 dB.

attenuates by 50 dB, the other boosts by 50 dB. Both stages add noise at –70 dB. The second stage will boost the noise of the first one. The input signal will leave this system with its level unaltered but with 50 dB extra noise.

To protect the signal from unnecessary boosts, the law of correct gain structure states:

Set the signal to optimum level as early as possible
in the signal chain, and keep it there.

To keep the signal at optimum level, we ideally want all the components it passes through to neither boost nor attenuate it; in other words, they should all be set to unity gain. Let us look at some practical examples:

- If the channel signal is too low, is it better to boost it using the input-gain control or the channel fader?
- If the input to a reverb unit is too low, should the master aux send boost it or the input gain on the reverb unit?
- If the overall mix level is too low, should we bring the channel faders up or the mix level?

The answers to all these questions are the first option in each case, these being the earlier stage in the signal chain. If we take the first question, for example, boosting the signal using the channel fader will also boost the noise added by the compressors and equalizers preceding it. Setting the optimum level at the input-gain stage means that the compressor, the equalizer, and the fader are fed with a signal at optimum level. If the channel fader is then used to attenuate the signal, it will also attenuate any noise that was added before it. If correct gain structure is exercised, the input gain of a reverb unit should be set to unity gain (0 dB), and the level to the reverb unit should be set from the master aux send.

Correct gain structure does not mean that signals should not be boosted or attenuated—it simply helps us to decide which control to use when there is more than one option. There are, of course, many cases where boosting or attenuating is the appropriate thing to do. For example, the makeup gain on a compressor, which essentially brings back the signal to an optimum level after the level loss caused by the compression process.

The digital console

Digital consoles might resemble the look of their analog counterparts, but they work in a completely different way. Analog signals are converted to digital as they enter the desk and are converted back to analog before leaving it. A computer—either a normal PC or purpose-built—handles all the audio in the digital domain, including processing and routing. The individual controls on the console surface are not part of any physical signal path—they only interact with the internal computer, reporting their state. The computer might instruct the controls to update their state when required (e.g., when a mix is recalled or when automation is enabled). Essentially, a digital desk is a marriage between a control surface and a computer.

The various controls can be assigned to many different facilities; for example, a rotary pot can control the signal level, pan position, aux send level, or equalizer gain, or even switch a compressor in and out. Most manufacturers build their consoles with relatively few controls and let the user assign them to different facilities at will. In addition, the user can select the channel to which various global controls are assigned. So, instead of having a set of EQ controls per channel strip, there will be only one global set, which affects the currently selected channel (Figure 10.33). While faders are motorized and buttons have an on/off indicator, pots are normally endless rotary controls and their position is displayed on a screen—from a small dot-display to a large, color touchscreen. The screen
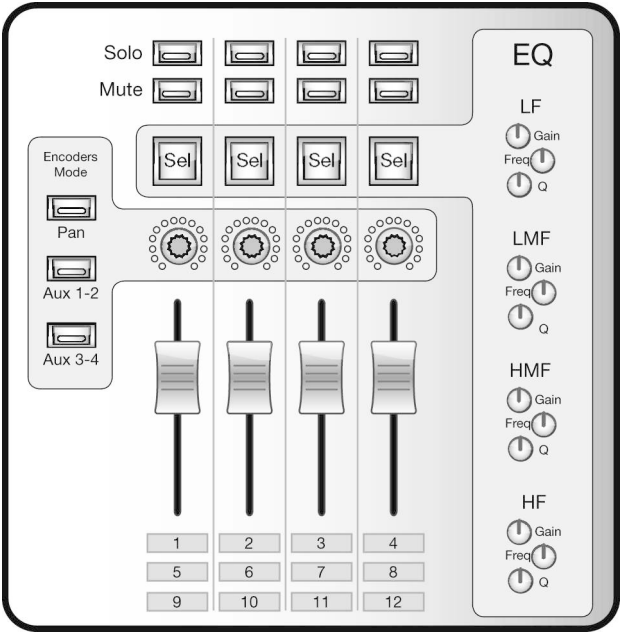
Figure 10.33 The layout of an imaginary mini-digital desk. Each strip has a single encoder, whose function is determined by the mode selection on the left. There is only one global set of EQ controls, which affects the currently selected channel.

also provides an interface for additional functions such as metering, scene store and recall, effect library management, and various configurations.

In this realm of assignable controls, each strip can control any channel. Essentially, one strip would suffice to control all the channels a desk offers, but this would be less convenient. Manufacturers provide a set of channel strips that control three-fold, fourfold, or other multiples of channels. For example, a desk with 16-channel strips might have 64 channels divided into four chunks (often called *layers* or *banks*). The user can select which layer is controlled by the channel strips at any given time (Figure 10.34). Effectively, all modern digital consoles are split consoles
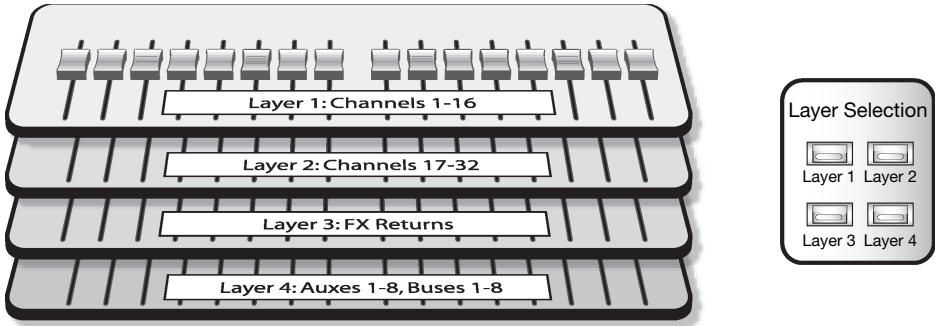
Figure 10.34 Layers on a digital desk. The 64 channels of this desk are subdivided into four layers, each represent 16 different channels. The layer selection determines which layer is displayed on the control surface.

where many strips are organized in a few layers. The strips can also be assigned to control master aux sends and group buses. (There is, however, a dedicated fader for the mix bus.) Just like software mixers, the differences between a channel, an auxiliary bus, and a group bus are blurred. With very few exceptions, the facilities available per channel are also available for the auxiliaries and groups—we can, for example, process a group with the same EQ or compressor available for a channel.

Assignable controls enable digital consoles to be substantially smaller than analog consoles but provide the same facilities. However, their operation is far less straight-forward than the what-you-see-is-what-you-get analog concept. Thus, not all digital consoles are indeed smaller. Leading manufacturers such as NEVE intentionally build consoles as big and with as many controls as their analog equivalents in order to provide the feel and speed of an analog experience. Still, these consoles provide the digital advantage of letting users assign different controls to different facilities. Users can also configure each session to have a specific number of channels, auxiliaries, or groups. The configuration of these large desks—which might involve quite some work—can be done at home, on a laptop, and then loaded onto the desk.

For the mixing engineer, digital desks aim to provide a complete mixing solution. Dynamic range processors and multiband equalizers are available per channel, very much like on large-format consoles. Most digital desks also offer built-in effect en-gines, reducing the need for external effects units (although these can be connected

Figure 10.35 The Yamaha O2R96 studio at SAE, London. This desk provides four layers, each consisting of 24 channels (or buses). Had this desk's layers been unfolded, it would be four times the size.

when needed in the standard way via aux sends). They let us store and recall our mixes (often called *scenes*) along with other configurations. A great advantage of digital consoles is that their automation engine allows us to automate virtually every aspect of the mix—from faders to the threshold of a gate. Perhaps the biggest design difference between digital and analog consoles has to do with insert points. Breaking the digital signal path, converting the insert send signal to analog, the insert return back to digital, and making sure everything is in sync requires far more complex design and often more converters. Many desks offer insert points that, instead of being located on the digital signal path, are placed before the A/D conversion. It is still possible to use external processors that way during mixdown but in a slightly different way from on an analog desk.

Console emulation

With analog consoles and outboard gear loved for their sound, and DAWs for their cost (and space) effectiveness, it doesn't take a genius to see the potential in some hybrid approach. While analog emulations have been around for a while (and are digital in nature), when these are at the center of how we work, we get a digital workflow with analog sound and to some extent "feel." The idea is simple—it's a DAW, but the processing is largely (if not wholly) done using analog emulation.

One example of such a hybrid approach is Universal Audio's Luna platform (Figure 10.36).—a DAW using dedicated hardware that allows both recording and mixing through analog emulation processors. To further support the analog workflow, Luna offers mixer extensions—processors that are integrated *per each channel*, with emulation of the pre-amp, EQ, or dynamics sections found on famous analog devices, as well as analog-summing emulation on buses. When used this way, both the interface and sound of the software mixer align with that of an analog console.

Another example is Harrison's MixBus32C DAW, the mixer of which integrates both the interface and sound of the analog 32C console.

The integration of specific emulation processors into software mixers provides analog-like look and sound in-the-box. The DAWs still support the any-of-many DAW workflow, where any plugin can be loaded onto any track; it is this universe of choice excess that so often hampers creativity when working on a DAW. Thus, sticking to the integrated processors these DAWs offer should translate into substantial reduction in choices and in turn a likely increase in speed and creativity—an example of enabling constraints. As a bonus we get an emulation of the beloved analog sound.
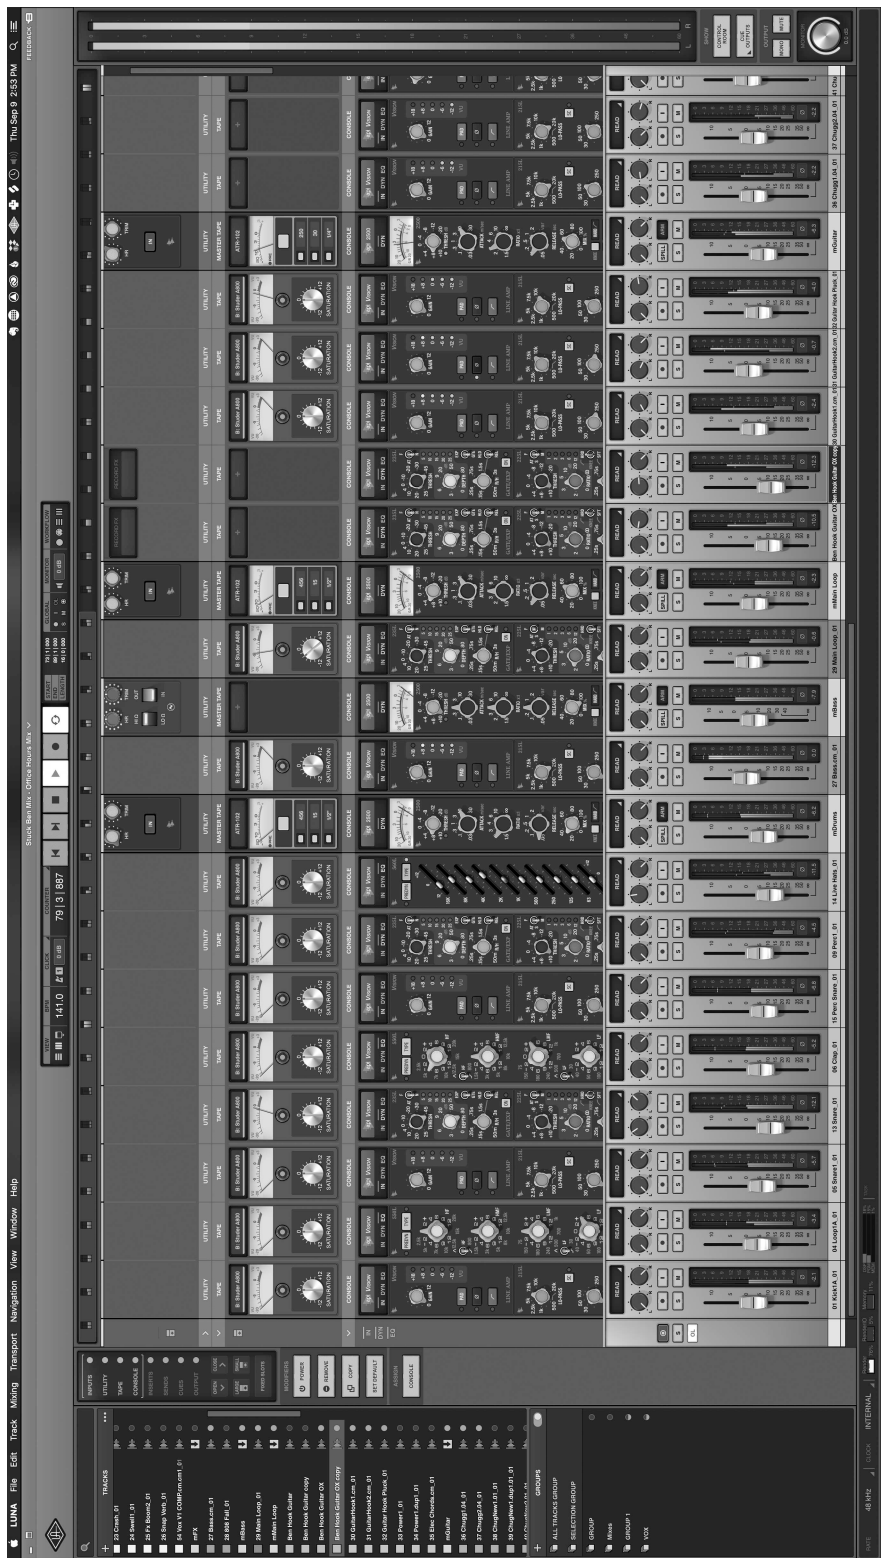
Figure 10.36 The Mixer window in Universal Audio's Luna DAW. The API Vision and Studer A800 emulation can be seen per audio channel as integrated extensions to the mixer board.

11 Software mixers

Computers have changed our lives, the music industry, audio engineering, and mixing. DAWs are gaining popularity, and as additional new and improved plugins are released each year, more and more professional mixing engineers are embracing software. While it would be fair to discuss all the audio sequencers currently available in the market, it would be impractical in the context of this book. So, with no disrespect to other similar products, the applications presented in this book are Steinberg's Cubase, MOTU's Digital Performer, Apple's Logic, and Avid's Pro Tools.

Audio sequencers let us mix *inside-the-box*. That is, they provide everything that is needed to complete a mix without external hardware (with the obvious exception of speakers). They can also be fully integrated with outboard gear and consoles when required. The software mixer provides the same core functionality as the mixing console: summing, processing, and routing. For routing, each software mixer offers a generous number of internal buses, used mainly as group and auxiliary buses. In audio-sequencers' jargon, these are simply called buses. All audio sequencers ship with processors and effects, either integrated into the software mixer or in the form of plugins that can be loaded dynamically. Third-party plugins extend the selection and can offer additional features and quality. All of the processing is done in the digital domain, with calculation carried out on the host CPU. Both summing and routing are relatively lightweight tasks; processors and effects are the main consumers of processing power, and the CPU speed essentially determines how many plugins can be used simultaneously in the mix. DSP expansions—either internal cards or external units—offer their own plugins that use dedicated hardware processors instead of the computer CPU, like Universal Audio's Apollo.

The physical inputs and outputs provided by the audio interface are available within the software mixer. If mixing were done wholly inside-the-box, a mixing station would only require a stereo output. Since we often bounce the final mix (rather than record it via the analog outputs), this stereo output might only be used for monitoring.

Tracks and mixer strips

Unlike a console-based setup, with audio sequencers there is no separation between the multitrack and the mixer—these are combined in one application. The multitrack is represented as a sequence window (or *arrangement, edit, project* window), where we

see the various tracks and audio regions. Whenever we create a new track, it appears in the sequence window, and a new **mixer strip** is created in the mixer window. The change in terminology is important—we lose the term channel, and talk about tracks and their associated *mixer strips* (Figure 11.1).

Tracks

Audio sequencers offer a few types of tracks. Those most likely to be used in mixing are:

• **Audio**—in the sequence window, these contain the raw tracks and their audio regions (references to audio files on the hard disk).
• **Aux**—used mainly for audio grouping and to accommodate effects as part of an aux send setup. In the sequence window, these tracks only display automation data.
• **Master**—most commonly represents the main stereo bus. In the sequence window, these tracks only display automation data.

Two more types of tracks, though less common, may still be used in mixing:

• **MIDI**—used to record, edit, and output MIDI information. Sometimes used during mixing to automate external digital units, or store and recall their state and presets. Might also be used as part of drum triggering.
• **Instrument**—tracks that contain MIDI data, which is converted to audio by a virtual instrument. On the mixer, these look and behave just like audio tracks. Might be used during drum triggering with a sampler as the virtual instrument.

Variations do exist. Cubase, for example, classifies its auxes as FX Channel and Group Channel.

Mixer strips

Figure 11.2 shows the audio mixer strips of various applications, and the collection of different sections and controls each one provides. Audio, aux, and instrument tracks look the same. Master tracks offer fewer facilities, while MIDI tracks are obviously different. We can see that, despite user interface variations, all the applications provide the same core facilities. The mixer strips are very similar to the channel strips found on a console—they have an input, and the signal flows through a specific path and is finally routed to an output. Figure 11.3 shows the simplified signal flow of a typical mixer strip. Since tracks can be either mono or stereo, mixer strips can be either mono or stereo throughout. However, a mixer strip might have mono input that changes to stereo along the signal flow (usually due to the insertion of a mono-to-stereo plugin).

The fader, mute button, and pan control need no introduction. The solo function will be explored shortly, while both automation and meters have dedicated chapters in this book. Other sections are:

• **Input selection**—determines what signal feeds the mixer strip. The options are either buses or the physical inputs of the audio interface. However, only aux strips truly work that way. Audio strips are always fed from the audio regions in the sequence window. When an audio track is recorded, the input selection
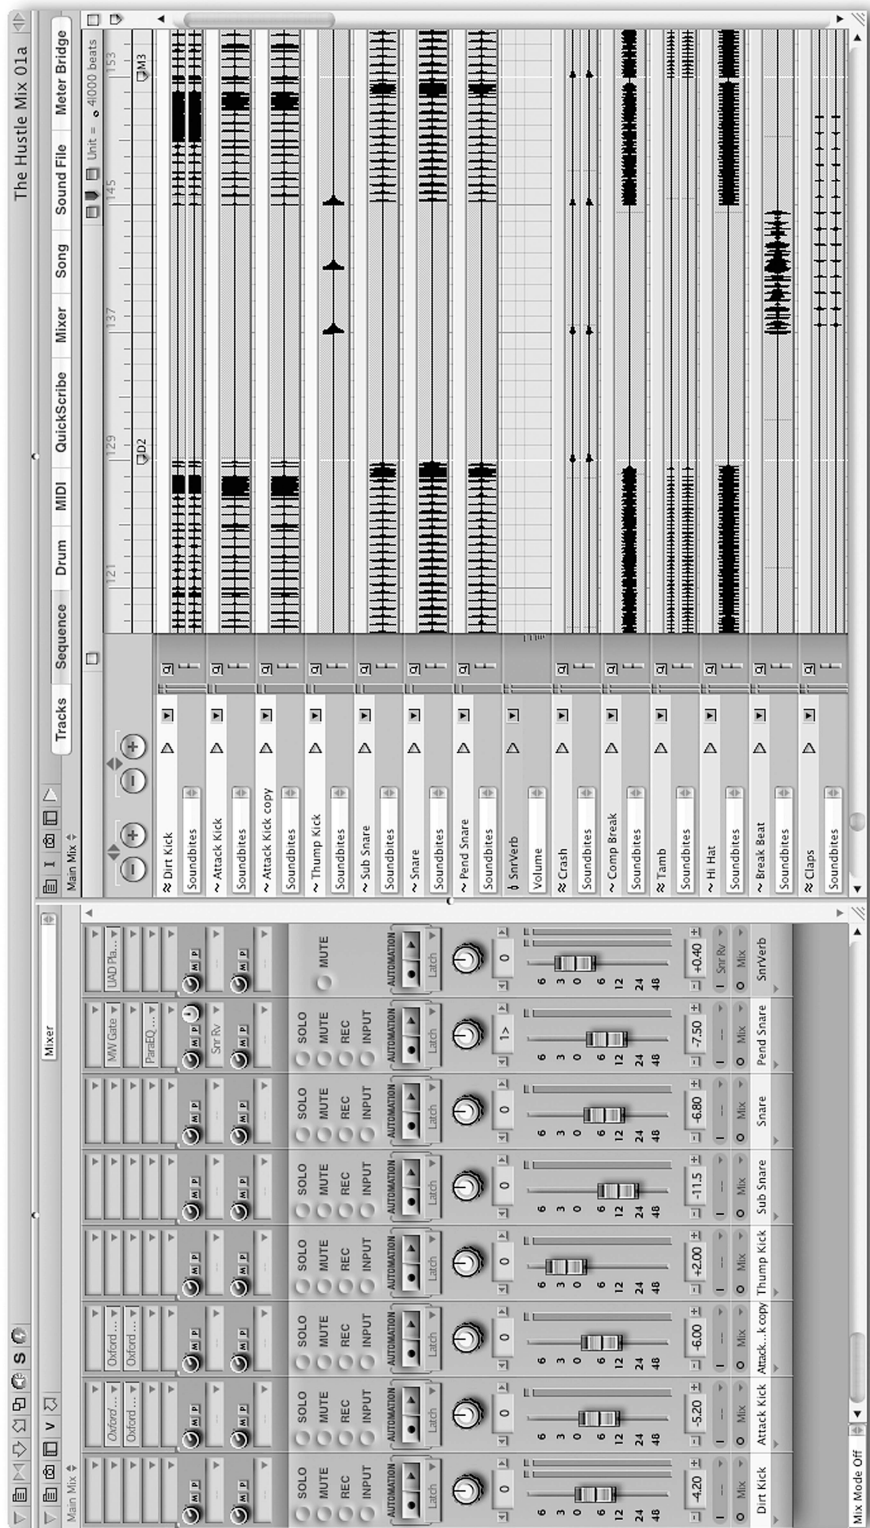
Figure 11.1 Digital Performer's mixer (left) and sequence window (right).
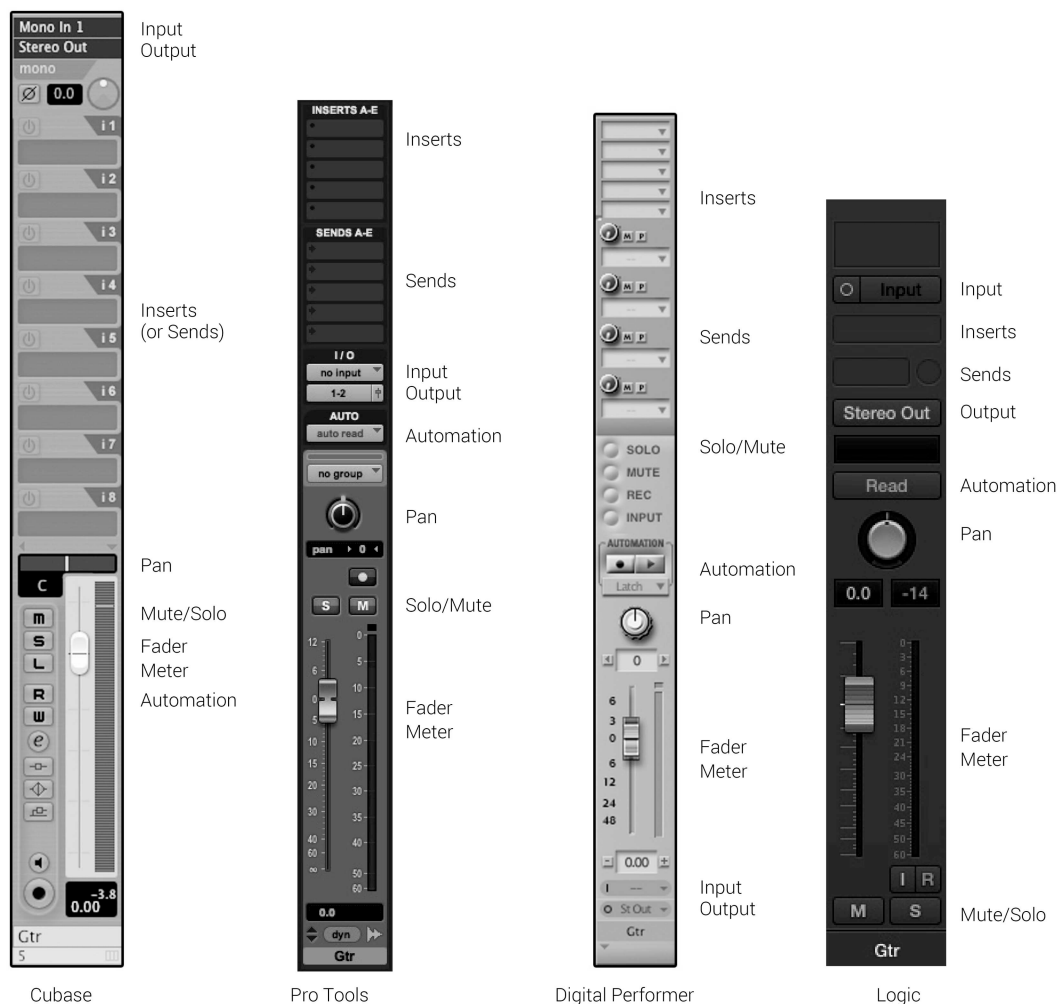
Figure 11.2 Audio mixer strips of different applications.

determines the source of the recorded material, but the audio is first stored in files, and only then fed to the mixer strip (therefore, any loaded processors are not captured onto disk).

- **Output selection**—routes the output of the mixer strip to either physical output or a bus. Some applications enable multiple selections so the same track can be routed to any number of destinations. This can be handy in very specific routing schemes (side-chain feed, for example). Commonly, mixer strips are routed to the main physical stereo output.
- **Insert slots**—lets us load plugins in series to the signal path, to facilitate the addition of processors. However, insert slots are also used to add effects, as we shall soon see. The order of processing is always top to bottom.
- **Send slots**—similar to the traditional aux sends, they feed either a pre- or post-fader copy of the signal to a bus, and let us control the send level or cut it altogether. Each send slot can be routed to a different bus, and the selection is made locally per mixer strip (i.e., two mixer strips can have their first slots routed to different buses).
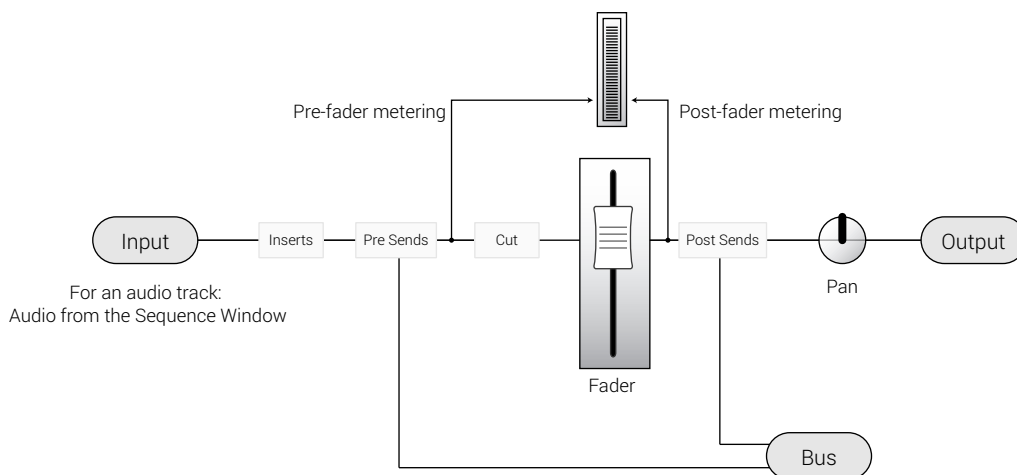
Figure 11.3 A signal flow diagram for a typical software mixer strip. Note that the single line can denote either mono or stereo signal path throughout—depending on the track.

Solos

Audio sequencers normally offer destructive in-place solo. Applications such as Cubase and Pro Tools HD also provide nondestructive solos. While some applications let us solo safe a track, others have a solo mechanism that automatically determines which channels should not be muted when a specific track is soloed. For example, soloing a track in Logic will not mute the bus track it is sent to.

Control grouping

Control grouping was mentioned in Chapter 10 as the basic ability to link a set of motorized faders on an analog desk. Audio sequencers provide much more evolved functionality than that. Each track can be assigned to one or more groups, and a group dialog lets us select the track properties we wish to link. These might include mute, level, pan, send levels, automation-related settings, and so on. Figure 11.4 shows the control grouping facility in Logic.

Control grouping, or just *grouping* as it is often termed, can be very useful with editing. However, when it comes to mixing, there is something of a love-hate relationship. On one hand, it removes the need for us to create an additional channel for an audio group. For example, if we have bass-DI and bass-mic, we often want the level of the two linked, and creating an audio group for just that purpose might be something of a hassle. On the other hand, control grouping does not allow collective processing over the grouped channels. There is very little we can do with control grouping that we cannot do with audio grouping. Perhaps the biggest drawback of control grouping is its actual essence—despite grouping a few tracks, often we may wish to alter some of them individually. Some applications provide more convenient means to momentarily disable a group or a specific control of a grouped track.

Figure 11.4 Logic's control grouping. The various backing-vocal tracks all belong to control group 1—called BV. The linked properties are ticked on the group settings window.

Routing

Audio grouping

Audio grouping provides the freedom to both individually and collectively process. As opposed to consoles, software mixers do not provide routing matrices. To audio-group tracks, we simply set the output of each included track to a bus, and then feed the bus into an auxiliary track. We can process each track individually using its inserts, and process the audio group using the inserts on the aux track. Figure 11.5 illustrates how this is done.

Sends and effects

As already mentioned, the addition of effects is commonly done using the send facility. Individual tracks are sent to a bus of choice, and, similarly to audio grouping, the bus is fed to an aux track. The effects plugin is loaded as an insert on the aux track, and we must make sure that its output is set to wet signal only. The aux track has its output set to the stereo mix. Figure 11.6 illustrates this.

Other routing

The input/output routing facility can be useful for more than grouping and sends. For example, a compressor loaded on the insert slots is pre-fader, but what if we wish to
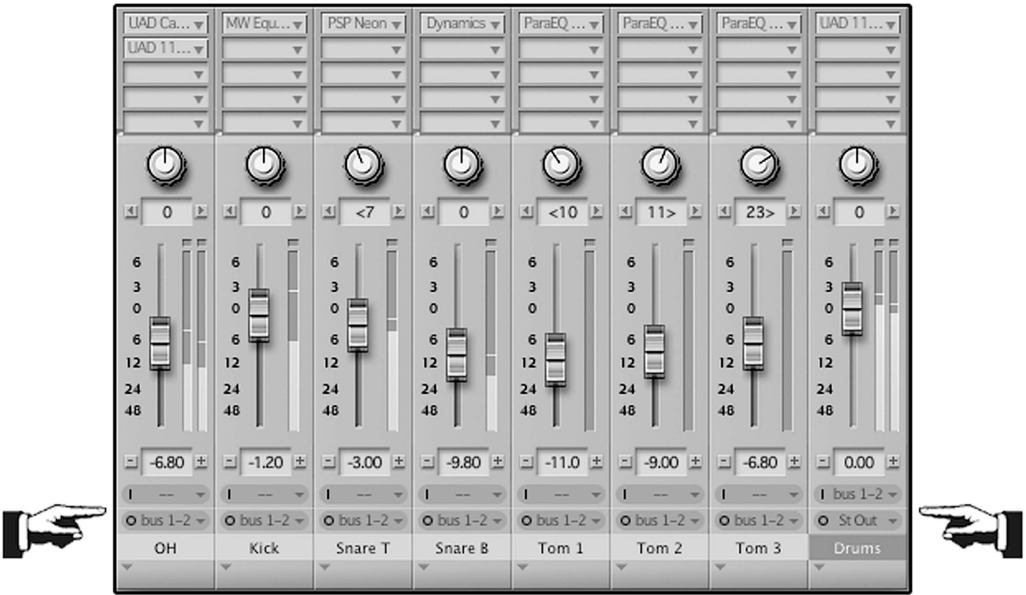
Figure 11.5 Audio grouping on a software mixer (Digital Performer). The outputs of the various drum tracks (OH to Tom 3) are routed to bus 1–2, which is selected as the input of an aux track (Drums); the aux track output is set to the stereo bus. By looking at the insert slots at the top, we can see that each track is processed individually, and there is also collective processing occurring on the drum group.

Figure 11.6 Sends on a software mixer (Pro Tools). Each vocal track is sent to bus 1–2. The aux track input is set to bus 1–2, and the output to the stereo bus. Also, the plugin wet/dry control is set to fully wet. In this setup, the audio tracks provide the dry signal (since each track has its output set to the stereo bus) and the wet signal is all the aux adds.
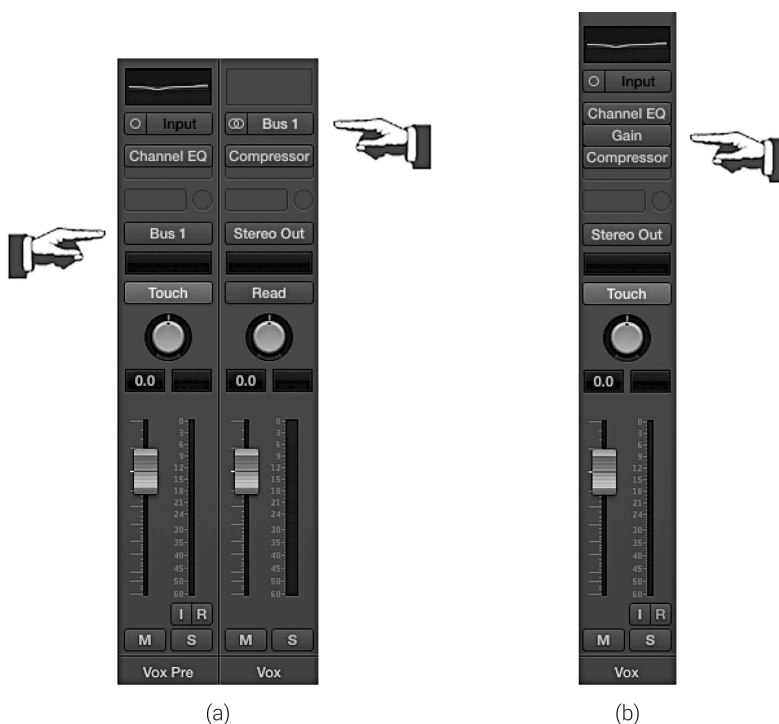
Figure 11.7 (a) Setting a track output to a bus and feeding it into an aux track lets us perform gain-riding before compression. The signal will travel through the EQ, the audio track fader (Vox Pre), which we gain-ride, the aux compressor, and the aux fader (Vox). (b) The same task can be achieved if we place a gain plugin before the compressor and automate the gain plugin instead of the fader. The signal will travel through the EQ; the gain plugin, which we automate; the compressor; and the audio track fader.

gain-ride the vocals and only then apply compression? We can route the output of the vocal track to an aux and insert the compressor on that aux. Figure 11.7a shows how this is done in Logic. A small trick provides a better alternative to this situation: we can insert a gain plugin before the compressor and gain-ride the former (Figure 11.7b). However, this might not be the ideal solution if we want to use a control surface to gain-ride the fader. Cubase provides an elegant solution to affairs of this kind—each mixer strip offers both pre- and post-fader inserts.

Naming buses (housekeeping tip)

All audio sequencers let us label the physical inputs, the outputs, and the buses. If we fully mix inside-the-box, we do not use physical inputs and we only use one stereo output, so the naming of these is less critical. But a complex mix might involve many buses, and very easily we can forget the purpose of each bus. Scrolling along the mixer trying to trace the function of the different buses is time consuming. This time can be saved if we label each bus as soon as it becomes a part of our mix. Figure 11.8 shows the same setup as in Figure 11.6, but with the reverb bus labeled.

Naming buses makes mixer navigation easier.

Figure 11.8 Pro Tools I/O setup (one place where buses can be labeled) and a mixer with a labeled bus.

The internal architecture

Integer notation

Digital waveforms are represented by a series of numbers. The sample rate defines the number of samples per second, and each sample holds a number that represents the amplitude of the waveform at a specific fraction of time. The possible values used to represent the amplitude are determined by the number of bits each sample consists of, along with the notation used to represent the numbers. If each sample is noted as a 16-bit integer, the possible values are 0 to 65,535 (negative integer notation is ignored throughout this section for simplicity). The highest amplitude a 16-bit sample can accommodate is represented by the number 65,535, which on a peak meter would light the full scale. Such a full-scale amplitude equals 0 dBFS, which is the highest possible level any digital system can accommodate. Mixing digital signals is done by summing (adding) sample values. Summing two samples with the value of 60,000 should result in 120,000. But since 16 bits cannot hold such a large number, a digital system will trim the result down to 65,535. Such trimming results in clipping distortion—normally, an unwanted addition. The same thing happens when we try to boost a signal beyond the highest possible value. Boosting by approximately 6 dB is done by doubling the sample value. Boosting a value of 40,000 by 6 dB should result in 80,000, but a 16-bit system will trim it down to 65,535.

Floating-point notation

Audio files are most commonly either 16- or 24-bit integer, and a D/A converter expects numbers in such form. However, audio sequencers handle digital audio with a different notation called *floating point*, which is slightly more complex than the integer notation. With floating-point, some of the bits (the *mantissa*) represent a whole number, while other bits (the *exponent*) dictate how this number is multiplied or divided. It might be easier to understand how the floating-point notation works if we define a simplified system where, in a four-digit number, the three rightmost digits represent a whole number (the mantissa), and the leftmost digit defines how many zeros should be added to the right of the mantissa. For example, with the value of 3,256, the whole number is 256 and three zeros are added to its right, resulting in 256,000. On the same basis, the value 0178 is equal to 178 (no added zeros). The most common floating-point notation, which has 23 bits of mantissa and 8 bits of exponent, is able to represent an enormous range of numbers, whether extremely small or extremely large. A 16-bit floating-point system supports much smaller and much larger values than a 16-bit integer system. As opposed to its integer counterpart, on a 16-bit floating-point system 60,000 + 60,000 does result in 120,000.

While the range of numbers supported by modern floating-point systems extends beyond any practical calculations humankind requires, there are some precision limitations attached. In the four-digit simplified system above, we could represent 256,000 and 178, but there is no way to represent the sum of both: 256,178. Floating point can support very small or large numbers, but no number can extend beyond the precision of the mantissa. A deeper exploration into the floating-point notation reveals that each mantissa always starts with binary 1, so this "implied 1" is omitted and replaced with one more meaningful binary digit. Thus, the precision of the mantissa is always one bit larger than the number of bits it consists of. For instance, a 23-bit mantissa has an effective precision of 24 bits.

The precision of the mantissa defines the dynamic range of a digital system, where each bit contributes around 6 dB (or precisely 20log2). Many people wrongly conclude that the famous 32-bit floating-point notation has 193 dB of dynamic range, when in practice the 24-bit mantissa only gives us around 144 dB. Two samples on such a system can represent a very high amplitude or an extremely low amplitude, which can be around 1,638 dB apart. However, when the two samples are mixed, the loud signals "win" the calculation, and any signal 144 dB below it is removed. In our four-digit simplified system, the result of 256,000 + 178 would be 256,000.

*There is some debate and misinformation (specifically on online forums) and lack of clarity when it comes to anything **64 bits**. This is largely because the computer architecture is often confused with the audio representation within a DAW, and a lack of understanding when it comes to floating-point audio.*

Old computers were using what is commonly referred to as "32-bit architecture." This means that the computer chips were designed to natively store, transfer, and process data that is 32 bits in length. A computer has memory and anything in the memory has its address, and because addresses were 32 bits, an application could only access 4 GB (2^{32}) of memory. As technology advanced to enable more than 4 GB of memory, this limitation

became a problem. Since the millennium, more and more computers and operating systems started to support the 64-bit architecture, and DAWs caught up over time. Nowadays, practically all new computers, operating systems, and DAWs are 64-bit ones.

This, however, has little to do in practice with the audio within a DAW. Even on 32-bit computers, plugin and DAW developers could (and some did) use 64-bit float audio processing; and even today on 64-bit machines, audio is still represented using 32-bit float. When it comes to audio, 32-bit float is often referred to as single precision*, and 64-bit float as* double precision*. Cubase offers an optional full 64-bit audio engine (meaning that throughout the system the audio is represented and processed using 64-bit float), although this only materializes if all the plugins used employ 64-bit processing. Other applications like Logic and Pro Tools use 64 bits for summing.*

What are the audio benefits in 64-bit float audio? It all has to do with dynamic range and distortion. But not the distortion caused by nonlinearity of processors like compressors; rather, it is the distortion added when two signals are summed and there are not enough digits to express the result of the summing (just like the 256,000 + 178 = 256,000 example above). In these cases, the most low-level data are being trimmed and distortion is added to the least significant bit of the audio, so it is very low in level and should be utterly negligible unless accumulated (by summing two or more distorted signals) or amplified.

The human ear has a dynamic range of 120 dB (which only applies in practice to a specific group of people in specific listening laboratories, yet it is the de-facto figure when discussing audio quality). The dynamic range of AD convertors hardly go above 120 dB, so regardless of the dynamic range of the audio, anything below –120 dB will simply not be heard.

32-bit float audio has a dynamic range of 144 dB, so we have a "floorroom" of 25 dBs—any distortion at the lower 25 dB range will not be heard. What does it take to generate more than 25 dB of distortion? In theory, summing 32 channels could yield 30 dB of distortion, but this depends on many assumptions that may not hold; for example, if the source audio is 16-bit integer (compared to 24-bit integer or 32-bit float), it would take many more channels before the distortion becomes audible. Yet, the main takeaway is that audible distortion within a 32-bit float audio engine is possible.

In comparison, 64-bit float audio offers 319 dB of dynamic range—so its floorroom is just shy of 200 dB. In theory, it would take the summation of more than 17 billion channels for the distortion to become audible.

So why not just use full-blown 64-bit audio engines? Because 64-bit engines take twice the memory of 32-bit ones. Intuitively, one would assume slower performance as there are also twice the bits to process; but on a modern 64-bit machine, arithmetic operations on 64-bit float numbers are only slightly slower, if not the same speed as that of 32-bit float ones. Ultimately, 64-bit audio boils down to a trade-off between memory and possible *distortion.*

The double-precision summing strategy seems like a good compromise. Audio is 32-bit float (less memory) but whenever signals are summed, they are summed into 64-bit float and then dithered back to 32-bit float. This process, which takes a bit more processing due to dithering, swaps distortion with (dither) noise; the latter accumulates far less than the former.

But then comes the question: How much extra memory is involved with a 64-bit float? A basic calculation is: buffer size * amount of channels * 2. So a project with a buffer size of

> *512 and 128 audio channels will consume 0.5 MB of memory with a 32-bit float, compared to 1 MB with a 64-bit float. Even if we take all into account and in practice the figure is 10-fold, we are still at the lower MB range—less than the memory consumed by a typical sampler plugin. 64-bit float throughout may not be such a bad idea after all.*
>
> *Often when there is no clear winner it is because it is too hard to call. This seems to be the case with the 32-bit float versus 64-bit float choice. There are many variables to account for and which is better may depend on the project and machine being used.*
>
> *Looking forward though, one could speculate that as years go by, the practical penalties in using 64-bit float engines will diminish and these will prevail. There are many benefits in a uniform audio representation, more so if machine native.*

What is in it for us?

The audio engine of audio sequencers is based on the 32- bit floating-point notation. The audio files that constitute the raw tracks are converted from integers to 32-bit float on-the-fly during playback. The audio data is kept as floating point throughout the mixer, and is only converted back to integers prior to the D/A conversion or when bouncing to an integer-based audio file (Figure 11.9). From the huge range of values 32-bit float numbers can accommodate, sample values within the application are in the range of –1.0 to 1.0 (the decimal point denotes a float number). This range, which can also be considered as –100 to 100 percent, was selected as it allows the uniform handling of different integer bit depths—a value of 255 in an 8-bit integer file, and a value of 65,535 in a 16-bit file are both full scale; thus both will be represented by 1.0 (100 percent).

We know already that a 16-bit integer system clips when two full-scale (65,535) samples are summed. Audio sequencers, with their 32-bit float implementation, have no problem summing two sample values of 1.0 (integer full scale)—software mixers live peacefully with sample values such as 2.0, 4.0, 808.0, and much larger numbers. With such an ability to overshoot the standard value range, we can theoretically sum a million tracks at integer full scale, or boost signals by around 764 dB, and still not clip on a 32-bit floating-point system. Practically speaking, even the most demanding mix would not cause clipping in the 32-bit floating-point domain.
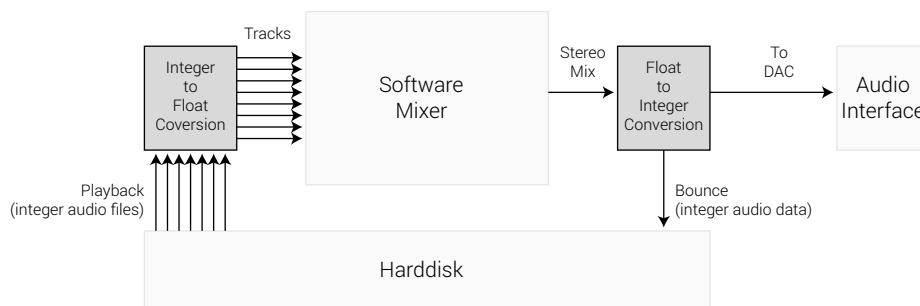
Figure 11.9 Macro audio flow in an audio sequencer and integer/float conversions.

Since 1.0 denotes integer full scale, audio sequencers are designed to meter 0 dB when such a level is reached. We say that 1.0 is the reference level of audio sequencers, and it denotes 0 dB. However, this is not 0 dBFS. Floating-point systems can accommodate values around 764 dB higher than 1.0, so essentially 0 dB denotes approximately −764 dBFS.

However, we are all aware that audio sequencers do clip, and that the resultant distortion can be quite obvious and nasty. The reason behind this is that at some point the float data is converted back to integer, and 1.0 is converted back to integer full scale (65,535 for 16 bits). During this conversion, a value above 1.0 is trimmed back to 1.0, which causes clipping distortion just like when integers are trimmed. The key point to remember here is that trimming—and its subsequent clipping distortion—can only happen during the float-to-integer conversion, nowhere else within the software mixer. This conversion is applied as the mix leaves the master track, therefore only overshooting signals at this stage will cause clipping distortion. While all mixer strips provide a clipping indicator with a threshold set to 1.0, none of these indicators denotes real clipping distortion—apart from the one on the master track.

Only clipping on the master track indicates clipping distortion.
All other individual tracks cannot clip.

As unreasonable as this fact might seem, it can be easily demonstrated by a simple experiment that is captured in Figure 11.11. We can set up a session with one audio and one master track, set both faders to unity gain (0 dB), then insert a signal generator on the audio track and configure it to generate a sine wave at 0 dB. If we then bring up the audio fader by, say, 12 dB, both clip indicators will light up (assuming post-fader metering), and clipping distortion will be heard. However, if we then bring down the master fader by 12 dB, the clipping distortion disappears, although the clipping indicator on the audio track is still lit up. By boosting the sine wave by 12 dB, it overshoots and thus the clipping indicator lights up. But, as the master fader attenuates it back to a valid level of 0 dB, it does not overshoot the master track, following which real clipping distortion occurs.

Track 11.1: Kick Source
The source kick track peaks at −0.4 dB.

Track 11.2: Kick Track and Master Clipping
By boosting the kick by 12 dB, it clips on the audio track. It also clips on the master track when its fader is set to 0 dB. The result is an evident distortion.

Track 11.3: Kick Track only Clipping
With an arrangement similar to that shown in Figure 11.10, the kick fader remains at +12 dB while the master fader is brought down to −12 dB. Although the post-fader clipping indicator on the audio track is lit, there is no distortion on the resulting bounce.

Figure 11.10 This Pro Tools screenshot shows a signal first being boosted by 12 dB by the audio track fader, and then attenuated by the same amount using the master fader. Despite the lit clipping indicator on the audio track, this setup will not generate any clipping distortion.

It is worth knowing that clipping-distortion will always occur when a clipping mix is bounced, but not always when it is monitored. This is dangerous since a supposedly clean mix might distort once bounced. Both Pro Tools and Cubase are immune to this risk as the monitored and bounced mixes are identical. However, when a mix clips within Logic or Digital Performer, clipping distortion might not be audible due to headroom provided by certain audio interfaces. Still, the mix will distort once bounced. Monitoring the master track clipping indicator in these applications is vital.

If individual tracks cannot clip, what is the point of having clipping indicators? A few possible issues make it in our interest to keep the audio below the clipping threshold. First, some processors assume that the highest possible sample value is 1.0—a gate, for example, will not act on any signals above this value, which represents the gate's highest possible threshold (0 dB). Second, the more signals there are crossing the clipping threshold, the more likely we are to cause clipping on the master track. While we will soon see that this issue can be easily resolved, why cure what can be prevented? Last, some plugins—for no apparent reason—trim signals above 1.0, which generates clipping distortion within the software mixer. To conclude, flashing clipping indicators on anything but the master track are unlikely to denote distortion, yet:

It is good practice to keep all signals below the clipping threshold.

Clipping indicators are also integrated into some plugins. It is possible, for example, that a boost on an EQ will cause clipping on the input of a succeeding compressor. It is also possible that the compressor will bring down the level of the clipping signal to below the clipping threshold, and thus the clipping indicator on the mixer strip will not light up. The same rule applies to clipping within the plugin chain—they rarely denote distortion, but it is wise to keep the signals below the clipping threshold.

Bouncing

It is worth explaining what happens when we bounce—whether it is intermediate bounces or the final mix bounce. If we bounce onto a 16-bit file, we lose 9 bits of dynamic range worth 54 dB. These 54 dB are lost forever, even if the bounced file is converted back to float during playback. Moreover, such bit reduction introduces either distortion or dither noise. While not as crucial for final mix bounces, using 16 bits for intermediate bounces simply impairs the audio quality. It would therefore be wise to bounce onto 24-bit files. Another issue is that converting from float to integer (but not the other way around) impairs the audio quality to a marginal extent. Some applications let us bounce onto 32-bit float files, which means that the bounced file is not subject to any degradation. However, this results in bigger files that might not be supported by some audio applications, including a few used by mastering engineers.

Bouncing to 16-bit files: bad idea, unless necessary.
24-bit: good idea.
32-bit float: excellent idea for intermediate bounces when disk space is not an issue.

Dither

Within an audio system, digital distortion can be the outcome of either processing or bit reduction. Digital processing is done using calculations involving sample values. Since memory variables within our computers have limited precision, sometimes a digital system cannot store the precise result of a specific calculation. For example, a certain fader position might yield a division of a sample value of 1.0 by 3, resulting in infinite 3s to the right of the decimal point. In such cases, the system might round off the result,

say to 0.333334. Since the rounded audio is not exactly what it should be, distortion is produced. Essentially, under these circumstances, the system becomes nonlinear, just like any analog system. Luckily, the distortion produced by calculations within our audio sequencer is very quiet and, in most cases, cannot be heard. The problems start when such distortion accumulates (a distortion of distortion of distortion, etc.). Some plugin developers employ double-precision processing. Essentially, calculations within the plugin are done using 64-bit float, rather than 32-bit float. Any generated distortion in such systems is introduced at far lower levels (around –300 dB). But in order to bring the audio back to the standard 32-bit float, which the application uses, the plugin has to perform bit reduction. Unless dither is applied during this bit reduction, a new distortion is produced.

The reason bit reduction produces distortion is that whenever we convert from a higher bit depth to a lower one, the system truncates the least meaningful bits. This introduces errors into the stream of sample values. To give a simplified explanation why: if truncating was to remove all the digits to the right of the decimal point in the sequence 0.1, 0.8, 0.9, 0.7, the result would be a sequence of zeros. In practice, all but the first number in our example are closer to 1 than to 0. Bit reduction using rounding would not help either, since this would produce rounding errors—a distortion sequence of 0.1, 0.2, 0.1, 0.3.

Dither is low-level random noise. It is not a simple noise, but one that is generated using probability theories. Dither makes any rounding or truncating errors completely random. By definition, distortion is correlated to the signal; by randomizing the errors, we decorrelate them from the signal and thus eliminate the distortion. This makes the system linear, but not without the penalty of added noise. Nonetheless, this noise is very low in level—unless accumulating.

These samples were produced using Cubase, which uses 32-bit float internally. A tom was sent to an algorithmic reverb. Algorithmic reverbs decay in level until they drop below the noise floor of a digital system. Within an audio sequencer, such reverb can decay from its original level by approximately 900 dB. The tom and its reverb were bounced onto a 16-bit file, once with dither and once without it.

Track 11.4: Rev Decay 16-bit
The bounced tom and its reverb, no dithering. The truncating distortion is not apparent as it is very low in level.

Track 11.5: Rev Tail 16-bit
This is a boosted version of the few final seconds of the previous track (the 16-bit bounced version). These few seconds are respective to the original reverb decay from around –60 to –96 dB. The distortion caused by truncating can be clearly heard.

Track 11.6: Rev Decay 16-bit Dithered
The bounced tom and its reverb, dithered version. Since the dither is at an approximate level of –84 dB, it cannot be discerned.

Track 11.7: Rev Tail 16-bit Dithered
The same few final seconds, only of the dithered version. There is no distortion on this track, but dither noise instead.

Track 11.8: Rev Decay 16-bit Limited
To demonstrate how the truncating distortion develops as the reverb decays in level, Track 11.4 is heavily boosted (around 60 dB) and then limited. As a result, low-level material gets progressively louder.

Track 11.9: Rev Decay 24-bit Limited
This is a demonstration that truncating distortion also happens when we bounce onto 24-bit files. The tom and its reverb are bounced onto a 24-bit file. The resultant file is then amplified by 130 dB and limited. The distortion error becomes apparent, although after a longer period compared to the 16-bit version.

Tom: Toontrack dfh Superior
Reverb: Universal Audio DreamVerb
Dither: t.c. electronic BrickWall Limiter

But the true power of dither is yet to be revealed: dither lets us keep the information stored in the truncated bits. The random noise stirs the very-low-level details into the bits that are not truncated, and thus the information is still audible even after bit reduction. The full explanation of how exactly this happens is down to the little-known fact that digital systems have unlimited dynamic range, but they do have a defined noise floor. This noise floor is determined, among other things, by the bit depth. In practice, a 16-bit system provides unlimited dynamic range, but it has its noise floor at around –96 dB. Human beings can discern information below that noise floor, but only to a point. If dither is not applied, the information in the truncated bits is lost, and distortion is produced. With dither, there is no distortion and the low-level details are maintained. However, at some point, they will be masked in our perception by the noise.

These samples were produced using Cubase. The original eight bars were attenuated by 70 dB, then bounced onto 16-bit files, once with dither and once without it. Then the bounced files were amplified by 70 dB.

Track 11.10: 8Bars Original
The original eight bars of drums and electric piano.

Track 11.11: 8Bars No Dither
The amplified no-dither bounce. Note how distorted everything is, and the short periods of silence caused by no musical information existing above the truncated bits.

Track 11.12: 8Bars Dither
The amplified dithered bounce. Note how material that was lost in the previous track can now be heard; for example, the full sustain of the electric piano. The hi-hat is completely masked by the dither noise.

Drums: Toontrack EZdrummer
Dither: t.c. electronic BrickWall Limiter

Theoretically, any processing done within a digital system can result in digital distortion. This includes compressors, equalizers, reverbs, and even faders and pan-pots. Due to its huge range of numbers, any floating-point system is far less prone to digital distortion compared with an integer system. Within an audio sequencer, most of this distortion is produced around the level of –140 dB, and for the most part would be inaudible. Manufacturers of large-format digital consoles apply dither frequently along the signal path. This is something audio sequencers do not do, but many plugins do (as part of double-precision processing). Dither has to be applied in the appropriate time and place. There is no point in us randomly adding dither, since doing so would not rectify the distortion—it would just add noise on top of it. The only instance where we should dither is as part of bit reduction. Such bit reduction takes place every time we bounce onto 16-bit files. We have just seen that bit reduction also causes distortion when bouncing onto 24-bit files, so there is a rationale for dithering even when bouncing onto such files. Only if we bounce onto 32-bit float files (Cubase, for example, enables this) do we not need to dither. The rule of dither is, therefore, very simple:

Apply dither once, as the very last process in the mix.

The generation of dither noise is not a simple affair. Different manufacturers have developed different algorithms; most of them involve a method called *noise-shaping* that lets us minimize the perceived noise level. Since our ears are less sensitive to noise at specific frequencies, shifting some of the dither energy to those frequencies makes the noise less apparent. There are different types of noise-shaping, and the choice is usually done by experimentation.

Most audio sequencers ship with dither capabilities. It might be an individual plugin, a global option, or part of the bounce dialog. In addition, more than a few plugins, notably limiters and loudness-maximizers, provide dither as well. Essentially, what the dither process does is add dither noise, and then remove any information below the destination bit depth. Figure 11.11 shows a dither plugin. The plugin is loaded on the master track in Pro Tools to establish itself as the very last process (even after the master fader). We can see the typical dither parameters: destination bit depth and noise-shaping type. In this case, the plugin is preparing for a bounce onto a 16-bit file.

Normalization

Figure 11.12 shows two audio regions. It is worth mentioning how the scale in Pro Tools works (or any other audio sequencer, for that matter). The middle mark is complete silence, otherwise roughly –151 dB; the top mark is 0 dB. Then every halving of the scale is worth –6 dB. The majority of the display shows the high-level part of the signal—there is as much space to show level variations between –12 and –18 dB as there is from –18 to –151 dB.

Consequently, it might seem that the right region is substantially lower in level than the left region, but in fact the left region is only 18 dB lower in level.

We often want our signals to be at optimum levels. An offline process called normalization lets us do this—it boosts the signal level to maximum without causing any clipping. For example, if the highest peak is at –18 dB, the signal will be boosted by 18 dB.
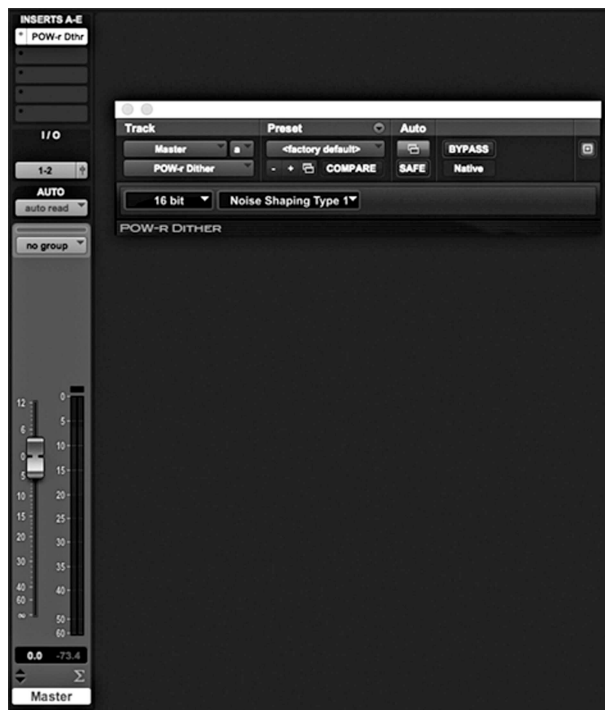
Figure 11.11 Pro Tools' *POWr Dither* plugin.
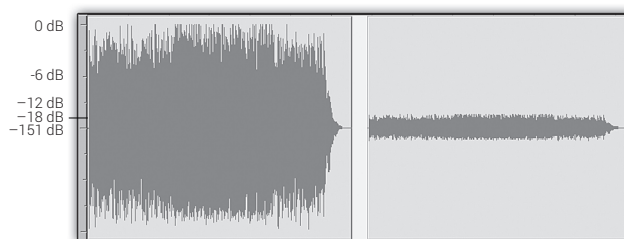
Figure 11.12 Waveform display scales. The left and right regions are identical, except that the right region is 18 dB quieter than the left one. The scales in editor windows are divided so that each halving results in –6 dB of level. The majority of what we see is the loud part of the waveform.

As an offline process, normalization creates a new file that contains the maximized audio. Normalization, however, has its shortcomings. During the process, rounding errors occur, resulting in the same distortion just discussed. This is especially an issue when 16-bit files are involved.

Normally, loud recordings are not placed next to quiet recordings as in Figure 11.12—we often simply have a whole track low in level. A more quality-aware approach to normalization involves boosting the audio in realtime within the 32-bit float domain. This can be easily done by loading a gain plugin on the first insert slot of the low-level track. This is very similar to the analog practice of raising the gain early in the signal-flow path when the signal on tape was recorded too low.

The master fader

Due to the nature of float systems, the master fader in audio sequencers is, in essence, a *scaling fader*. It is so-called since all it does is scale the mix output into a desired range of values. If the mix clips on the master track, the master fader can scale it down to an appropriate level. It is worth remembering that a master track might clip even if none of the individual channels overshoots the clipping threshold. In these situations, there is no need to bring down all of the channel faders—simply use the master fader. In the opposite situation, where the mix level is too low, the master fader should be used to scale it up.

This is also worth keeping in mind when bouncing. The master fader (or more wisely a gain plugin on the master track) should be used to set the bounced levels to optimum. Determining the optimum level is easy if we have a numeral peak display—we have to reset the master track's peak display and play the bounce range once. The new reading on the peak display will be equal to the amount of boost or attenuation required to bring the signal to 0 dB (or to the safety –3 dB) (Figure 11.13).

The master track in Pro Tools is worth mentioning. As opposed to all the other track types, a master track in Pro Tools has its inserts post-fader. This is in order to enable dithering at the very last stage in the digital signal chain. The consequence of this is that any dynamic processing applied on the master track, such as limiting, will be affected by the master fader. If we automate, or fade out the master fader, the amount of limiting and its sonic effect will vary (Figure 11.14a). Although seldom desired, such an effect can be heard on a few commercial mixes. To fix this issue, individual tracks are
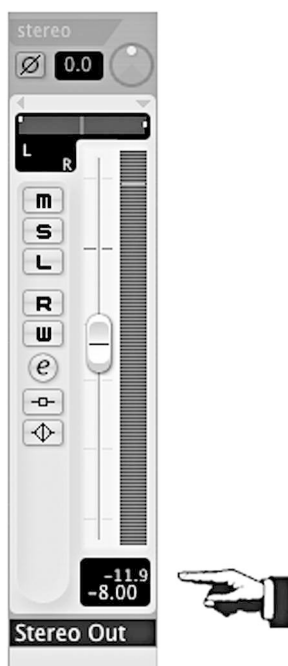
Figure 11.13 The master track in Cubase after the bounced range has been played. The top value shows the peak measurement (–11.9 dB) and tells us that the master fader can be raised by nearly 11.9 dB before clipping will occur.
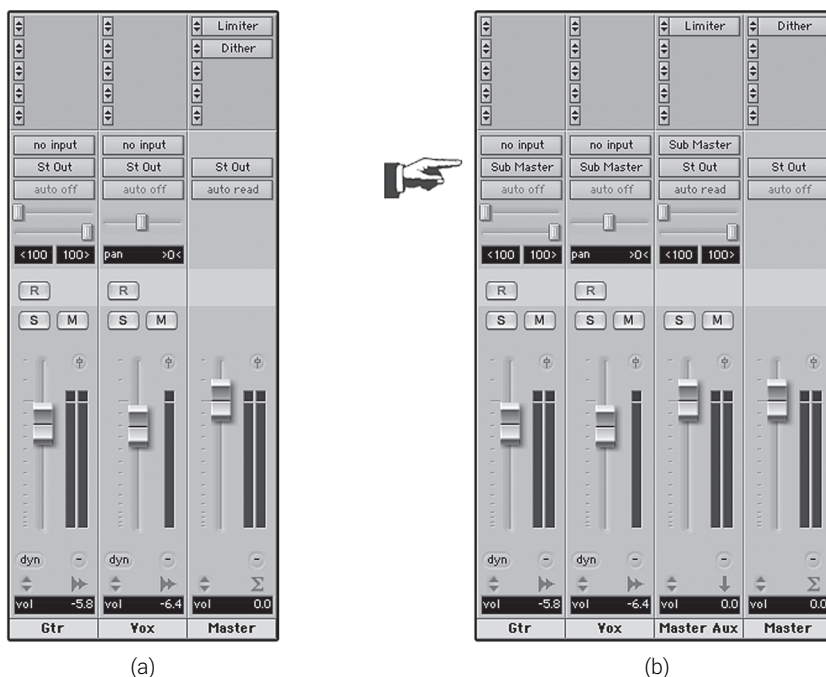
Figure 11.14 (a) Since the limiter on the master track is post-fader, any fader movements will result in varying limiting effect. (b) The individual audio tracks are first routed to an auxiliary, where limiting is applied, and only then sent to the main output, where dithering takes place post-fader. Note that in this setting both the aux and the master faders succeed the limiter, so both can be automated. Automating the aux is preferred, so the master fader remains as an automation-independent scaling control.

routed to a bus that feeds an auxiliary, and the limiting is applied on the auxiliary. The dithering is applied on the master track (Figure 11.14b). Here again, having both pre- and post-inserts per mixer strip, like in Cubase, would be very handy.

The playback buffer

Audio sequencers handle audio in iterations. During each iteration, the application reads samples from the various audio files, processes the audio using the mixer facilities and plugins, sums all the different tracks, and delivers the mix to the audio interface for playback. Each iteration involves a multitude of buffers; each represents a mono track or a mono bus. The playback buffer size determines how many audio samples each of these buffers holds. For example, if the playback buffer size is 480 samples, during each iteration 480 samples from each file are loaded into the track buffers; each plugin only processes 480 samples, and the different tracks are summed into the mix-bus buffer, which also consists of 480 samples. Once the mix buffer has been calculated, it is sent to the audio interface. If the buffer size is 480 samples and the playback sample rate is set to 48,000, it will take the audio interface exactly 10 ms to play each mix buffer. While the audio interface is playing the buffer, another iteration is already taking place. Each iteration must be completed, and a new mix-buffer delivered to the audio interface, before the previous buffer finishes playing (less than 10 ms in our case). If an iteration takes longer than that, the audio interface will have nothing to play, the continuous
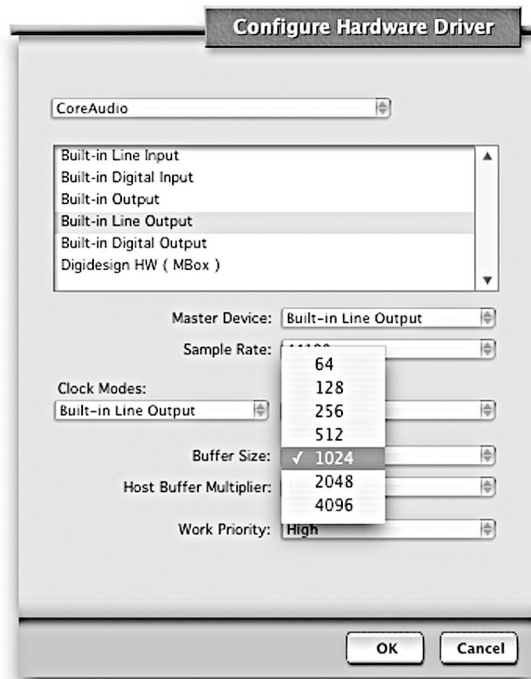
Figure 11.15 Digital Performer's configuration window where the playback buffer size is set.

playback will stop, and a message might pop with a CPU overload warning. A computer must be powerful enough to process audio in real time faster than it takes to play it back.

The playback buffer size determines the monitoring latency of input signals, and the smaller the buffer size, the lower the input latency. This is important during recording, but less so during mixdown, where input signals are rarely involved. Smaller settings, however, also mean more iterations—each with its own processing overhead—so fewer plugins can be included in the mix. With a very small buffer size (say, 32 samples), some applications might become unpredictable, start producing clicks, or skip playback. However, depending on the application, a large buffer size can affect the accuracy of some mix aspects, such as automation. When choosing the buffer size, it should be large enough to allow a sufficient number of plugins and predictable performance; anything larger than that might not be beneficial. A sensible buffer size to start experimenting with is 1,024 samples.

Plugin delay compensation

We have just seen that plugins process a set number of samples per iteration. Ideally, a plugin should be given an audio buffer, process it, and output the processed buffer straight away. However, some plugins can only output the processed buffer a while after, which means their output is delayed. There are two reasons for such plugin delay. The first involves plugins that run on a DSP expansion such as the UAD card. Plugins of this kind send the audio to the dedicated hardware, which returns the audio to the

plugin after it has been processed. It takes time to send, process, and return the audio; therefore, delay is introduced. The second reason involves algorithms that require more samples than those provided by each playback buffer. For example, a linear-phase EQ might require around 12,000 samples in order to apply its algorithm, which falls short of the common 1,024 samples the plugin gets per iteration. To overcome this, the plugin accumulates incoming buffers until these constitute a sufficient number of samples. Only then, after some time has passed, does the plugin start processing the audio and output it back to the application.

Plugin delay, in its least harmful form, causes timing mismatches between various instruments, as if the performers are out of time with one another. A mix like that, especially one with subtle timing differences, can feel very wrong, although it can be hard to pinpoint timing as the cause. In its most harmful form, plugin delay causes comb filtering, tonal deficiencies, and timbre coloration. For example, if plugin delay is introduced on a kick track, it might cause comb filtering when mixed with the overheads.

Luckily, all audio sequencers provide automatic plugin delay compensation. In simple terms, this is achieved by making sure that all audio buffers are delayed by the same number of samples.

12 Phase

What is phase?

Phase, in sound theory at least, describes the time relationship between two or more sine waves of the same frequency. Phase is measured in degrees. A sine wave starts when its cycle starts—denoted by 0°. One complete cycle signifies 360° (which is also 0°), half a cycle 180°, a quarter of a cycle 90°, and so forth. If one sine wave reaches a quarter of its cycle when a cycle begins on another sine wave of the same frequency, the two are said to be 90° out of phase (Figure 12.1).
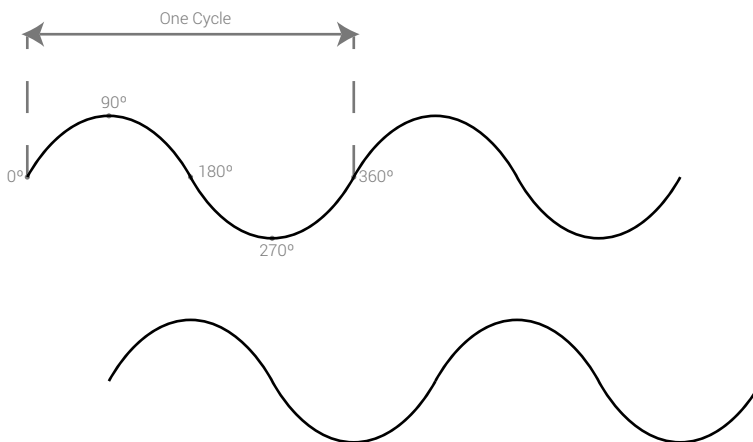
Figure 12.1 Two identical sine waves 90° out of phase with each other.

The degree of phase shift is dependent on the time gap between the two sine waves, but also on their frequency. For the same time gap, different frequencies will have different phase shifts. Complex waveforms, like all recorded and synthesized waveforms (apart from sine waves), contain many frequencies. If two complex waveforms are delayed with respect to one another, each frequency within them will have its own degree of phase shift.

In mixing, we aren't quite interested in sine waves, but rather in complex waveforms. When two **similar** waveforms are mixed, their phase interaction is of interest. We have to define *similar* first:

- **Identical waveforms**—as the name suggests, these are similar in every way and usually the outcome of duplication. For example, by duplicating a vocal track in an audio sequencer, we get two identical waveforms. By sending a snare to two different groups, we also create two duplicate signals.
- **Waveforms of the same event**—two microphones capturing the same musical event. For example, two microphones fronting an acoustic guitar, or a kick-mic and the overheads—although the kick-mic captures mostly the kick, and the overheads capture the whole drum kit, both recordings will have the same kick performance in them.

It is also important to define what is dissimilar: two takes of any instrument, no matter how similar in performance, do not count as similar waveforms. Based on this, double-tracking is not subject to phase consideration.

> *There is one exception to the above. A kick's sound is largely determined by its first cycle, which can be similar enough between two kicks—even if these were recorded using a different drum and microphone. In practice, phase mostly becomes a factor when we layer two kick samples.*

There are three types of phase relationships between similar waveforms:

- **In phase** or **phase-coherent**—when the waveforms start at exactly the same time.
- **Out of phase** or **phase shifted**—when the waveforms start at different times. For example, if we delay or nudge the duplicate of a vocal track, but not the original. Also, if two microphones are placed at different distances from an instrument, they will capture the same waveform, but the farthest microphone will do so slightly later.
- **Phase inverted**—when both waveforms start at exactly the same time, but their amplitude is inverted (i.e., mirrored across the zero line). Many people confuse phase inversion with a 180° phase shift, but the two are not the same. When two waveforms other than sine waves are 180° out of phase, they do not start at the same time and rarely does their amplitude appear inverted (Figure 12.2).
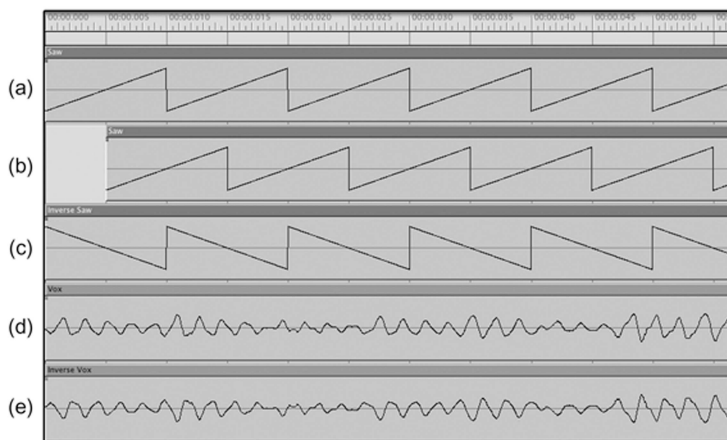
Figure 12.2 Five mono tracks in Digital Performer. (a) The original rising sawtooth. (b) The same sawtooth only 180° out of phase with the original. (c) The original sawtooth phase inverted. (d) A vocal track. (e) The same vocal track phase inverted.

One thing that the Haas effect tells us is that an audible delay will only be perceived if the time gap between two similar waveforms is bigger than 30–40 ms. In this text, 35 ms will be used, and we will regard the *Haas window* as any time gap between 1 and 35 ms. We normally consider phase when time gaps fall into the Haas window. In mixing terms, a time gap longer than 35 ms is an audible delay and phase is no longer of interest.

Problems

Phase problems occur when we sum similar waveforms that are either phase shifted or phase inverted. If the time gap involved is smaller than 35 ms, one should expect **comb filtering** and its subsequent tonal alterations and timbre coloration. If two waveforms are phase inverted with one another, they will **cancel** each other and cause level attenuation; if both are of equal amplitude, they will cancel each other completely and result in complete silence.

It should be said that all of this holds true when the two waveforms are summed *internally*—either electronically or mathematically within the mixer. If each of the waveforms is played through a different speaker, the result is different, as we shall soon see. However, if the stereo output is folded to mono, the same phase problems will arise.

Phase problems on recorded material

Raw tracks, to our great misfortune, often involve some phase issues—especially those of recorded productions. One example is a guitar amplifier recording made with a microphone a few feet from the grill. With both the amplifier and the microphone close to a reflective floor, the direct and the later reflected sound are likely to produce noticeable comb filtering when they hit the microphone. In cases such as this, comb filtering is imprinted into a single track, and usually there is little the mixing engineer can do. However, when two or more tracks represent the *same take* of the *same instrument*, some enhancements can be made. Here are a few usual suspects for tracks that might benefit from phase correction:

- **Top/bottom** or **front/back tracks**—snares are often recorded with two microphones in a top and bottom arrangement, while kicks and guitar cabinets might be recorded with two, front and back, microphones. Since the two microphones are positioned on opposite sides of the sound source, they are likely to capture opposite sound pressures (contrary to belief, this is not always the case, especially if the microphones are at different distances from the instrument). Recording engineers are aware of this and usually invert the phase of one microphone before it is recorded. This way, the two tracks arrive ready for mixing in phase.
- **Close-mic** and **overheads**—it is possible to have the kick, snare, or any other close-mic either phase inverted or phase shifted with the overheads. Again, a good recording engineer will take the overheads as a reference and make sure all drum tracks are phase coherent with the overheads.
- **Mic** and **direct**—in addition to the conventional microphone recording, guitars (bass, in particular) might also be recorded using a DI box, or direct output from the

amplifier. Being kept in the electronic domain, direct signals travel at almost the speed of light. Sound takes some extra time to travel through the air between the amplifier and the microphone, resulting in a slightly delayed instance of the mic track compared with the direct track. This is one example of phase shifting that, once corrected, can result in dramatic timbre improvements.

These tracks demonstrate the effect of phase shift on the timbre of a bass guitar. Such phase shifts can be the outcome of mixing direct and mic tracks, or two microphones that were positioned at different distances from the source:

Track 12.1: Bass In Phase
The two bass tracks perfectly in phase.

Track 12.2: Bass 1 ms Shift
A 1 ms time gap between the two tracks results in loss of high frequencies, definition, and attack. (It is worth mentioning that this is the most primitive form of a digital low-pass filter—by summing a sound with a very short delayed version of itself, high frequencies are softened while low frequencies are hardly affected.)

Track 12.3: Bass 2 ms Shift
With a 2 ms time gap, the loss of low frequencies becomes evident. The track sounds hollow and powerless.

Track 12.4: Bass 5 ms Shift
A 5 ms time gap produces a severe coloration of timbre, and a distant, weak sound. These two tracks show the effect of phase inverting one of two snare tracks:

Track 12.5: Snare Top and Bot
Both the snare top and bottom are mixed as recorded.

Track 12.6: Snare Top with Bottom Inverted
With the snare bottom track inverted, some snare body is lost. The effect is not always similar to the one heard here, or as profound.

Phase problems are not always audibly obvious. In the case of phase inversions, we go through the usual suspects and see what happens when we engage the phase-invert switch. Sometimes the difference is noticeable, whether for good or bad; sometimes it is not. We can also zoom into associated tracks to look for phase inversions. Zooming can also be useful when hunting for phase shifts. To correct phase shifts, we nudge the late audio region backward in time until it is aligned with its associated track. We can also delay the early track until it is aligned with the late track, but this less straight-forward solution is mostly used when nudging is not available, as with a tape multi-track. When aligning two phase-shifted tracks, it is not always possible to determine the exact nudging amount by eye, and different values result in different timbres. Fine phase corrections are mostly done by ear.

> *Phase checks and corrections are better carried out before*
> *mixing onset. Associated tracks are checked for phase*
> *inversions and phase shifts.*

Figure 12.3 Phase shift on guitar recordings. This screenshot shows two guitar recordings that were captured simultaneously using microphones at different distances from the amplifier. The time gap, approximately 5 ms, is evident here.

Phase problems during mixdown

Although rare, phase problems can also be created during mixdown; here are a few examples how:

- **Lack of plugin delay compensation**—uncompensated plugin delay is an open invitation for phase problems. For example, if a compressor plugin delays the snare by 500 samples, chances are that when mixed with the overheads comb filtering will occur.
- **Digital-to-analog conversions on outboard gear**—digital units convert the audio from analog to digital and back, which takes some time. This introduces a short delay, but one long enough to cause comb filtering if the signal is then summed with an identical track.
- **Very short delays**—short delays that fall into the Haas window might cause comb filtering if the original and the delayed sounds are summed internally.
- **Equalizers**—as part of their design, equalizers cause group delay for a specific range of frequencies. Phase problems are especially noticeable with large gain boosts.

Tricks

We've mentioned that when two identical waveforms are not in phase, but each is played through a different speaker, the result is quite different from comb filtering. Two mixing tricks are based on such a stereo setup. With both, two identical mono signals are sent to a different extreme, and one of the signals is either delayed or phase-inverted. To distinguish the two, we will call the unaltered signal the *original signal* and the copy, which is either delayed or phase-inverted, the *ghost copy*.

The Haas trick

The Haas trick was not invented by Helmut Haas, but it is essentially a demonstration of the Haas effect. Haas was interested in what happens when an initial sound is quickly succeeded by similar sounds from various directions. His findings teach us that the directivity of the sound is determined solely by the initial sound provided that (a) the successive sounds arrive within 1–35 ms of the initial sound; and (b) the successive sounds are less than 10 dB louder than the initial sound. Although the successive sounds do not give any directivity cues, they still play a spatial role. The Haas trick simply involves an original signal panned to one extreme, and a ghost copy, which is delayed by 1–35 ms, sent to the other (Figure 12.4).

The Haas trick is usually achieved in one of two ways. The first involves panning a mono track hard to one channel, duplicating it, and then panning the duplicate hard to the opposite channel and nudging the duplicate by a few milliseconds (Figure 12.5). The second involves loading a stereo delay on a mono track, and setting one channel to have no delay and the other to have a short delay between 1 and 35 ms.

The Haas trick results in a wide, open, spacious sound; although the sound can be localized to the non-delayed channel, there is some sense of unfocused stereo image. It can be applied during mixdown for three main purposes:

- **To fatten sounds panned to the extremes**—using the Haas trick on instruments already panned to the extremes can make them sound bigger, especially in a crossed arrangement. For example, it is common to double-track distorted guitars and pan each mono track to an opposite channel. Applying the Haas trick on each guitar track (sending the delayed duplicate to the opposite channel as in Figure 12.6) results in a fatter, more powerful effect.
- **As a panning alternative**—sometimes when panning a mono track, all the panning options seem less than ideal. For example, in a sparse arrangement of three mono tracks—vocal, bass, and guitar—chances are that both the bass and the vocal will be panned center. Panning the guitar to the center will place it in the busiest,
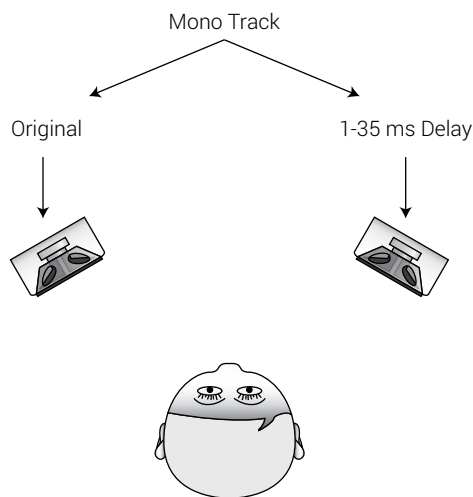
Figure 12.4 The Haas trick. A mono track is sent to both left and right, but one channel is delayed by 1–35 ms.
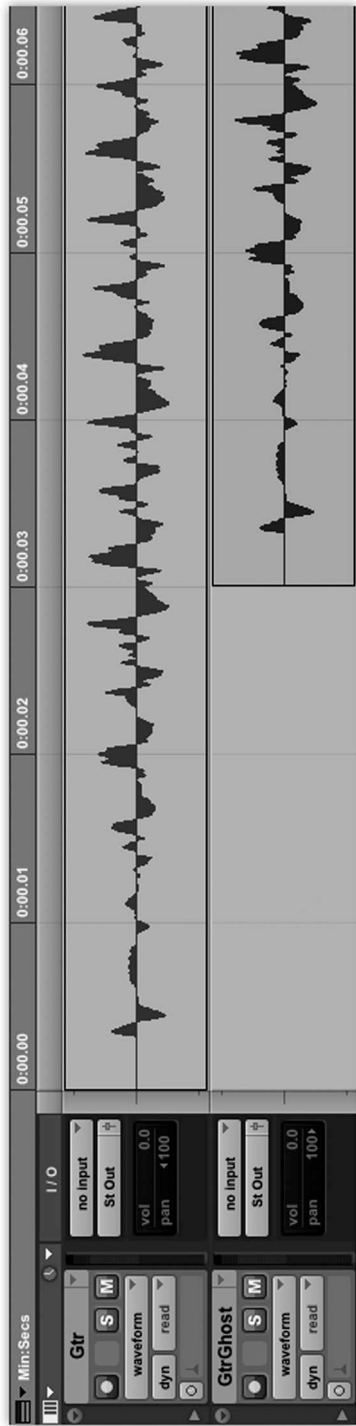
Figure 12.5 The Haas trick using duplicates. One way to achieve the Haas trick is by duplicating a mono track, nudging the duplicate a few milliseconds in time (30 ms in the screenshot above), then panning each track to an opposite extreme.

Gtr 1 Original
+
Gtr 2 Ghost

Gtr 2 Original
+
Gtr 1 Ghost
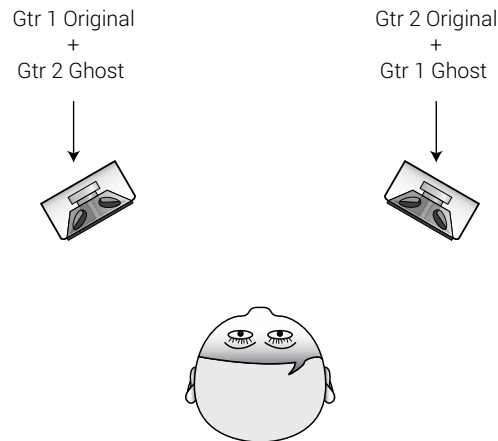
Figure 12.6 The Haas trick in crossed setup. Guitar 1 is sent to the left and its delayed copy to the right. Guitar 2 is sent to the right and its delayed copy to the left.

high-masking area of the mix (and will also create a monophonic mix). Panning the mono guitar to one side or the other will cause stereo imbalance. By applying the Haas trick, we can open up the monophonic guitar sound, give it some stereo width, and place the guitar in a low-masking area.

- **More realistic panning**—the ear uses amplitude, time, and frequency differences for the localization of sounds. A pan-pot works by only altering amplitude; therefore, the results are less natural. The Haas trick adds to the standard panning method time differences, and with a filter on the delayed channel we can also tuck on some frequency differences. However, this application is limited to instruments already panned to one extreme—we cannot pan the original signal anywhere else.

Track 12.7: dGuitars Stereo
Two distorted guitars, each panned hard to a different extreme, before applying the Haas trick.

Track 12.8: dGuitars Haas Crossed
The same two guitars after applying the Haas trick in crossed arrangement, as shown in Figure 12.6. The delay times are set to around 28 ms.

Plugin: Digidesign DigiRack Mod Delay II.

One of the settings we can control with the Haas trick is the **amount of delay** applied on the ghost copy. Different delay times give slightly different effects, so it is up to the ear to decide which is the best effect. One thing to consider is what happens when the mix is folded to mono. To be sure, the resultant comb filtering will have its impact (often loss of high frequencies), but some delay times, mostly longer ones, sum to mono more gracefully than others, so while checking in mono we look for the least-destructive delay time.

Another setting we can control is the **level** of the ghost copy. The Haas effect only applies if the ghost copy is less than 10 dB louder than the original signal. Depending on the delay time, it might not apply with figures even smaller than 10 dB. Assuming that both the original and the ghost copy are initially at the same level, boosting the level of the ghost copy will appear at first to balance the stereo image. With slightly more boost, a rather confusing, unnatural, yet interesting effect is achieved—it can be described as an audible delay where it is impossible to say which extreme is the delayed one.

Track 12.9: dGtr Original
The original distortion guitar in mono.

Track 12.10: dGtr Haas 5 ms
The Haas trick applied on this track with the right channel delayed by 5 ms.

Track 12.11: dGtr Haas 30 ms
The Haas trick applied on this track with the right channel delayed by 30 ms. Compared to the previous track, the effect here appears wider and fuller.

Track 12.12: dGtr Haas 5 ms Mono
Summing Track 12.10 (5 ms delay) to mono results in very evident comb-filtering effect.

Track 12.13: dGtr Haas 30 ms Mono
There is still coloration of the sound when Track 12.11 (30 ms delay) is summed to mono, but this coloration is neither as evident nor as obstructive as the effect with 5 ms delay.

Plugin: Digidesign DigiRack Mod Delay II.

Attenuating the ghost copy below the level of the original signal makes the effect less noticeable and the stereo image of the original signal somewhat less vague. Altogether, it is worth experimenting with different levels, and making the final decision based on taste and feel.

The following tracks demonstrate how the relative level of the ghost copy affects our perception of the Haas trick. When applied, the delay time of the ghost copy was set to 23 ms.

Track 12.14: Guitar Src
The source track used in the following samples. This is the mono recording with no delay involved.

Track 12.15: Guitar Panned
In the following tracks, the source guitar is panned hard-left and the ghost copy hard-right. This track, which is given for reference purposes, only involves the source track panned hard-left, without the ghost copy.

Track 12.16: Guitar Haas 9 dB Down
The Haas effect is applied on this track with the ghost copy 9 dB below the original. The ghost copy on the right channel is not easily discerned. But comparing this track to the previous one would reveal the effect. In the previous track, the guitar is distinctly isolated on the left channel. In this track, the image still clearly appears to come from the left channel, but in a slightly more natural way. This is an example of how the Haas trick can yield more realistic panning and subtly richer sound. Yet, even in a sparse mix, the ghost copy on the right channel could be masked to an extent that would conceal the effect; therefore, it might need to be louder.

Track 12.17: Guitar Haas 6 dB Down
In this track, the ghost copy is 6 dB below the original. It is only slightly more noticeable. The guitar image still clearly appears to come from the left channel.

Track 12.18: Guitar Haas 3 dB Down
With the ghost copy 3 dB below the original, it becomes more noticeable, and the resultant effect is richer. The guitar image still appears to come from the left channel.

Track 12.19: Guitar Haas 0
Both the original and the ghost copy are at the same level in this track. As Haas suggested, the sound image still appears to come from the left channel.

Track 12.20: Guitar Haas 3 dB Up
The ghost copy is 3 dB above the original. The effect is becoming confusing, although the image still appears to come from the left channel. It is interesting that the image should still come from the left channel if you move your head toward the right speaker.

Track 12.21: Guitar Haas 6 dB Up
With the ghost copy 6 dB above the original, the effect is becoming very confusing indeed. Most listeners would still perceive the image as if coming from the left, but there is some sense of balance between the left and right speakers. One way or another, the overall effect is wide and it creates an unfocused image (which can be appealing if used in the right context).

Plugin: Digidesign DigiRack Mod Delay II.

The Haas trick can be enhanced in two ways. First, applying a **filter** on the ghost channel—usually a low-pass filter—can bring about more natural results, and might even reduce comb filtering when the mix is folded to mono. This is especially useful when the Haas trick is used to achieve more realistic panning. When the Haas effect is applied to fatten sounds, another enhancement is possible. While using a delay on the ghost channel, we **modulate** the delay time with low rate and low depth (Figure 12.7). This gives a richer impact and bigger size to the effect. While the results are not always suitable, it is worth experimenting.

The Haas trick is often only expedient if used on one, maybe two instruments in the mix. When used on more, it can clutter the extremes and result in an overall unfocused stereo image—doing more harm than good to the mix. As with many other mixing tricks, overdoing means loss of impact.
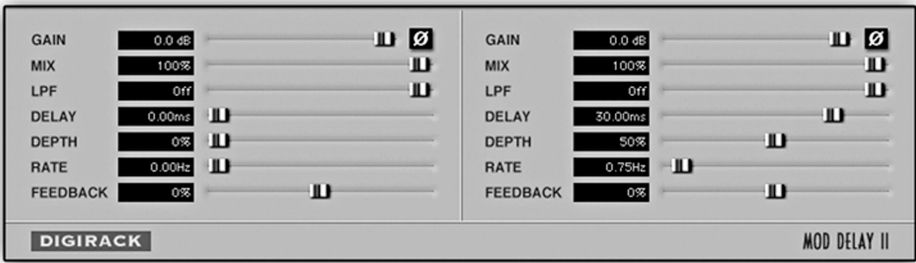
Figure 12.7 The modulated Haas trick using a delay plugin (Digidesign *DigiRack Mod Delay II*). This mono-to-stereo instance of the plugin involves a left channel that is neutralized with 0 ms delay time. The right channel is delayed by 30 ms and modulated using 50 percent depth and 0.75 Hz rate.

Track 12.22: dGtr Haas 30 ms Modulated
The same setup as in Track 12.11, only in this track the delay time is modulated using 50 percent depth and 0.75 Hz rate. The exact setup can be seen in Figure 12.7. Compared to Track 12.11, the effect here is richer and provides a slight sensation of movement.

Plugin: Digidesign DigiRack Mod Delay II.

The out-of-speakers trick

With the Haas trick, we delay the ghost copy. With the out-of-speakers trick, we keep it time-aligned to the original signal, but we invert its phase (Figure 12.8). This means that the sound arriving at one ear is phase-inverted with the sound arriving at the other. In nature, this happens (for specific frequencies only) when the sound source is located to one side of the head. The two phase-inverted signals emitted from each speaker first travel to the near ear, but shortly after both arrive at the far ear (Figure 12.9).
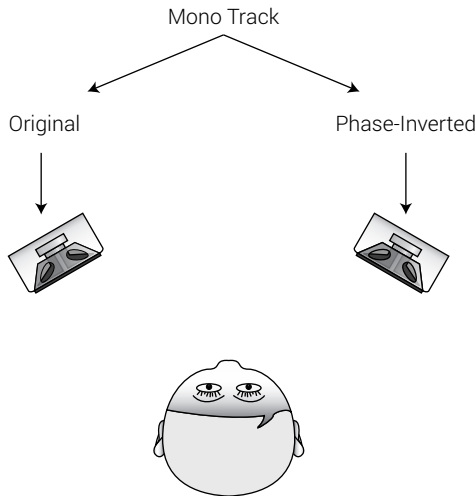
Figure 12.8 The out-of-speakers trick. A mono track is sent to both left and right, but one channel is phase-inverted.
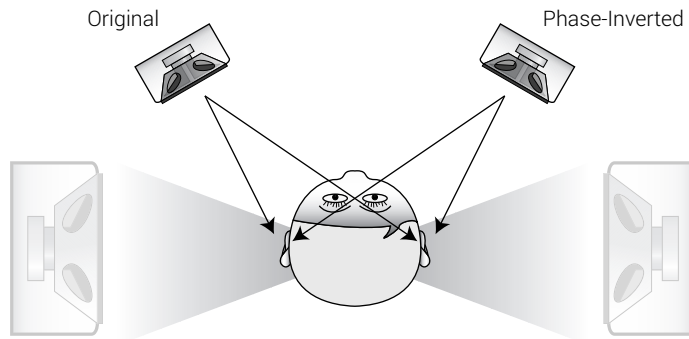
Figure 12.9 With the out-of-speakers trick, the original and inverted signals will first reach the ear closer to the speaker from which each signal is emitted, and shortly after the other ear. This creates the impression that the sound is arriving from wide side angles.

The outcome of this is sound that seems to arrive at both sides of the head at the same time. Since different frequencies are localized slightly more forward or backward, the final effect is of sound coming from around you, rather than simply from your left and right.

When heard for the first time, the out-of-speakers trick can make some people over-whelmingly excited—not only does sound appear to arrive from outside the speakers, it also seems to surround you in a hallucinatory fashion. To add to the excitement, an instrument on which the out-of-speakers trick is applied will disappear completely when the mix is folded to mono (provided that the original and the ghost copy are at exactly the same level, as they usually are with this trick).

Track 12.23: 8Bars both Mono
Both the drums and electric piano are monophonic and panned center.

Track 12.24: 8Bars ePiano OOS
The out-of-speakers trick is only applied to the electric piano. When listening in mono, the piano should disappear.

Track 12.25: 8Bars Drums OOS
The out-of-speakers trick is only applied to the drums. When listening in mono, the drums should disappear.

Track 12.26: 8Bars both OOS
The out-of-speakers trick is applied to both instruments. Listening to this track in mono should yield silence.

Plugin: Logic Gain.
Drums: Toontrack EZdrummer.

The trick, however, is only in full effect when the listener's head is on the central plane between the two speakers. Also, our ears only localize sounds coming from the sides

based on low-frequency content. The wavelength of high frequencies is too short compared with the distance between our ears. High frequencies can easily change their phase a few times while traveling the distance between our ears. Scientifically speaking, 1 kHz is roughly the frequency below which side localization is experienced. Consequently, it is mostly frequencies below 1 kHz that give the impression of surrounding the head. The out-of-speakers trick is therefore mostly effective with low-frequency sounds.

> *The out-of-speakers trick causes low-frequency material to appear as if*
> *coming from all around us. Instruments to which the trick is*
> *applied disappear in mono.*

These tracks demonstrate the varying effect of the out-of-speakers trick in relation to frequencies. An arpeggiated synth line involves a band-pass filter that is swept up, then down. The kick and hi-hats provide a center-image reference, on which the trick is not applied.

Track 12.27: Arp Mono
The arpeggiated line in mono, panned center.

Track 12.28: Arp Out of Speakers
The arpeggiated line with the out-of-speakers trick applied. It is possible to hear that the higher the band-pass filter is swept, the narrower the synth image becomes.

Track 12.29: Arp Toggle
The out-of-speakers effect is toggled every two beats. The greater effect at lower frequencies should be easier to discern here.

Plugin: Logic Gain.

A sound that appears to surround your head in a hallucinatory fashion is clearly the definition of anti-focus. We can apply the out-of-speakers trick on an instrument, but we should not expect the instrument to relate well to the sound stage we are building. What's more, the instrument will disappear in mono, so it would not be wise to apply this trick on important instruments, say the beat of a hip-hop track. The out-of-speakers trick is used as a special effect, a sonic gimmick, and is usually applied on the least-important instruments, or ones that appear for a very short time in the mix.

> *The out-of-speaker trick is used as a special effect, occasionally,*
> *and mostly on the least-important tracks.*

To conclude this section, it should be noted that the out-of-speakers trick is not strictly vinyl-proof. In fact, the worst-case scenario for vinyl cutting is phase-inverted low-frequency content on the left and right channels. For this reason, mastering engineers sum the low end of the mix to mono, which cancels out, to some degree, any instruments to which this trick has been applied. If the mix is to be cut on vinyl, mixing engineers may be required to submit a vinyl edit from which this trick is omitted.

13 Faders

Even before the invention of the multitrack, various microphone sources were balanced using faders and summed before being recorded. The summing devices used in early recording studios had faders all right, but these were rather bulky rotary knobs, often 3.5″ in diameter. When trying to balance more than two microphones, hands had to jump between knobs. One engineer to foresee the advantage of employing linear faders was the legendary Tom Dowd. By building a control section with small linear faders, he could control each fader using a finger, and so balance more than two microphones at a time. Technically speaking, a fader is a device that can fade sound, whether it is rotary or linear. Today, the term *fader* is widely associated with linear controls, while *pots* denote rotary control. This book follows this convention.

Faders are the most straightforward tools in the mixing arsenal. They are the main tools for coarse-level adjustments, while both equalizers and compressors are used for more sophisticated or fine-level control. Faders are the first thing we approach when starting to mix, and sometimes part of the final fine-level adjustments. It might seem unnecessary to say anything about faders—you slide them up, the level goes up; slide them down, the level drops. But even faders involve some science worth knowing. They are not as straightforward as one may think.

Types

Sliding potentiometer

The simplest analog fader is based on a sliding potentiometer. The amplitude of analog signals is represented by voltage, and resistance is used to drop it. Inside a fader, there is a resistive track on which a conductive wiper slides as the knob is moved. Different positions along the track yield different resistance and thus different degrees of level attenuation. One shortfall of the sliding potentiometer is that it cannot boost the level of audio signals passing through it (however, the potentiometer can be made to behave as if boosting by placing a fixed-gain amplifier after it). While the actual circuitry involved is more complex, we can consider this type of fader as one into which an audio signal enters and leaves, as shown schematically in Figure 13.1.
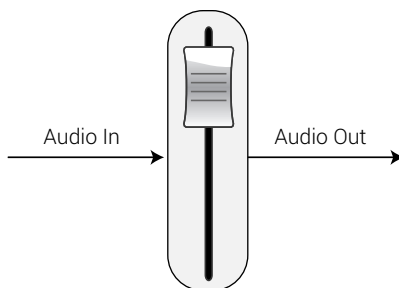
Figure 13.1 A schematic illustration of a sliding potentiometer.
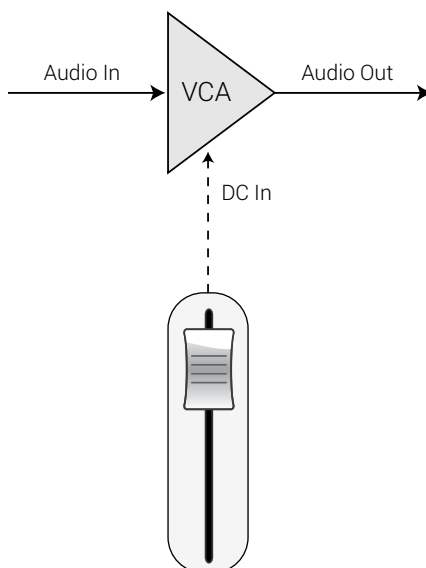
Figure 13.2 A VCA fader. The fader only controls the DC voltage sent to a VCA. The VCA is the component that boosts or attenuates the audio signal as it passes through it.

VCA fader

A VCA fader is a marriage between a voltage-controlled amplifier (VCA) and a fader. A VCA is an active amplifier through which audio passes. The amount of boost or attenuation applied on the signal is determined by incoming DC voltage. The fader, through which no audio flows, only controls the DC voltage sent to the amplifier. Figure 13.2 demonstrates this.

One advantage of the VCA concept is that many DC sources can be summed before feeding the VCA. A channel strip on an analog console can be designed so many level-related functions are achieved using a single VCA. This shortens the signal path and reduces the number of components in it, resulting in a better noise performance. SSL consoles, for example, are designed around this concept—DC to the VCA arrives from the VCA fader, VCA groups, cut switch, and automation computer. Figure 13.3 illustrates this.
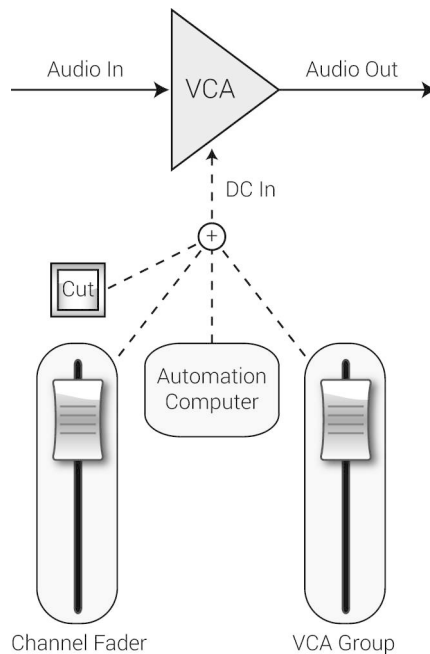
Figure 13.3 A single VCA can be used with many level-related functions.

Digital fader

A digital fader simply determines a coefficient value by which samples are multiplied. For example, doubling the sample value—a coefficient of 2—results in approximately 6 dB boost. Halving the sample value—a coefficient of 0.5—results in approximately 6 dB attenuation. Other coefficient values can be easily calculated based on a simple formula.

Scales

Figure 13.4 shows a typical fader scale. There are perhaps as many fader scales as manufacturers and software developers. While the actual numbers might be different, the principles discussed in this section are common to the majority of faders.

The typical scale unit is dB, which bears a strong relationship to the way our ears perceive loudness. Most often, the various steps are based on either 10 dB (which is subjectively a doubling or halving of perceived loudness) or 6 dB (which is approximately doubling or halving of voltage or sample value). The 0 dB point is also known as *unity gain*—a position at which a fader neither boosts nor attenuates the signal level. Any position above 0 dB denotes a level boost, while anything below denotes attenuation. Most faders can boost; therefore, they provide an *extra-gain*. Common extra-gain figures are 6, 10, and 12 dB. We should regard this extra-gain as an emergency range—using it implies non-optimal gain structure and possible (often minor) degradations to the sound. Ideally, this extra gain should not be used, except on rare occasions when all the mix levels are set and a specific track still calls for a bit of push above unity gain. At the very bottom of the scale, there is –∞ (minus infinity), a position
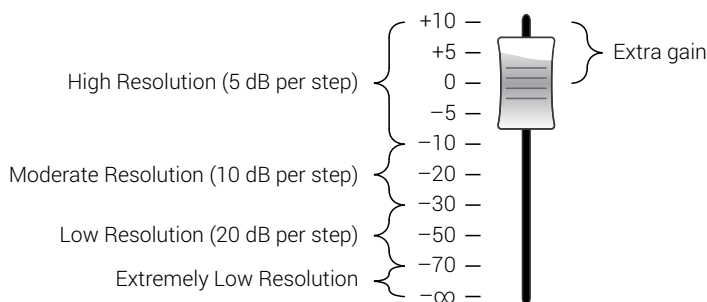
Figure 13.4 A typical fader scale.
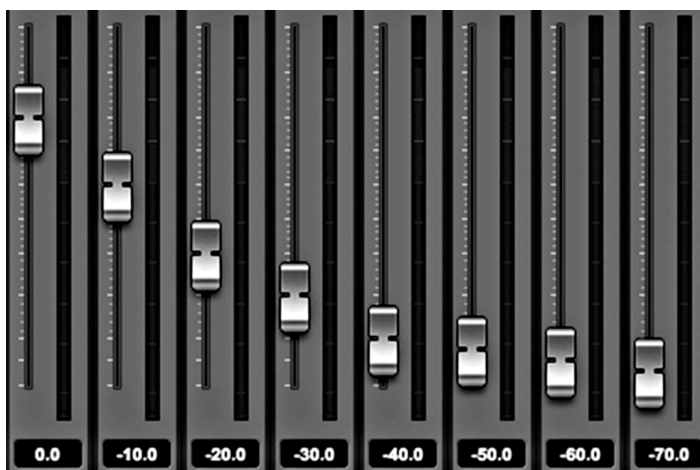
Figure 13.5 A demonstration of the uneven fader scale. These eight Pro Tools faders are set at –10 dB intervals. This demonstrates how hard it can be to set any value within the –70 to –60 dB range, while it is quite easy within the –20 to 0 dB range. Modifier keys can accompany mouse movements for more precise results.

at which the signal is inaudible (often achieved by attenuation of 120 dB or more, but some analog desks cut the signal when the fader reaches this position).

One very important thing that becomes evident from Figure 13.4 is that, although the different scale steps are evenly distributed, the dB gaps between them are not consistent. At the top of the scale (between –10 dB and +10 dB), the steps are 5 dB apart; at the bottom (–30 to –70 dB), the steps are 20 dB apart. This means that sliding the fader at its lower end will cause more drastic level changes than sliding it at its higher end. Another way to look at this is that at the top end of the fader level, the changes are more precise—clearly something we want while mixing. This suggests that we want to have our faders in the high-resolution area of the scale; provided we do not use any extra gain, the area between –20 and 0 dB is the most critical. One useful function found on some audio sequencers is the ability to key in the required level or have more precise control over the fader position using modifier keys.

Another consequence of the uneven scale is that, when we gain-ride a fader, the lower it goes, the more drastic the attenuation becomes (and the opposite effect for raising). If we fade out a track, for example, we might have to slow down as we approach the bottom of the scale.

Working with faders

Level planning

A common question probably asked by every mixing engineer at least once is: So which fader goes up first and where exactly does it go? There are a few things we have to consider in order to answer this question.

Faders like to go up. It should not come as a surprise if throughout a mixing session one by one all the faders go up—a few times—ending up at exactly the same relative positions. This is the outcome of the louder-perceived-better axiom. Here is an example of how easily things can go wrong; you listen to your mix and find yourself unsure about the level of the snare. You then boost the level of the snare. This is likely to make it more defined, so the move appears to be right. A few moments after, when working on the vocals, you realize that they do not stand out since the snare is nearly as loud, so you boost the vocals. The vocals should now stand out all right. Then, you find that the kick seems too weak since the snare is way louder. So you boost the kick. But now you are missing some drum ambiance so you boost the overhead, then the toms, then the cymbals, then the bass, then the guitars, and before you know it you are back to square one—you are unsure about the snare again.

Faders like to go up.

There are a few ways to solve this cyclical syndrome. First, understand that faders are the least-sophisticated tools for making something stand out. Second, listening in mix-perspective minimizes the likelihood of individual level moves. Third, it takes some discipline to stick to the plan, especially when it comes to level boosts. After absorbing these three ideas, you might still find that faders like to go up, so the ultimate solution is this:

Leave some extra gain available.

The extra-gain dead-end is shown in Figure 13.6. It happens when a fader is fully up, but the instrument can still use some gain. This scenario can be solved in various ways, but it would be much easier planning the levels right in the first place, so the extra-gain dead-end never happens. We have said already that the fader's extra gain is better kept for an emergency. So we normally do not want any faders above 0 dB at the early mixing stage. If we also take into account that faders like to go up, it might be wise to leave even more additional gain, so there is still some distance to go before we start using the extra gain. For example, it might be wise to start the mix with the loudest instrument set to –6 dB (which gives us 6 dB of virtual extra gain on top of the standard extra gain).

Level planning requires setting the loudest instrument of the mix first, and then the rest of the faders in relation to it. An example of the opening mix moves for a production where the

Figure 13.6 The extra-gain dead-end. The vocal track is at the top position of +12 dB, which means that no additional level boost can be applied by the fader if needed.

lead vocal is expected to be the loudest would involve setting the lead vocal track to –6 dB and then setting the overheads level with respect to the lead vocal. The vocal track can then be muted if one wishes to start mixing from the drums. Rough mixes can help as well—if the rough mix ends with the highest fader at +4 dB, that fader should be at 0 dB (or slightly below it) in the real mix, with the rest of the levels set in comparison to this level. Another strategy is to create a quick balance mix (faders only) of key tracks before mixing onset. Then adjust the fader levels so the highest fader is on or just below 0 dB.

> *Set initial levels in relation to what you think will be the loudest*
> *instrument in the mix.*

So that is the digital version of level planning. Analog mixing works slightly differently, and we can borrow the analog wisdom for the benefit of digital. Unlike software sequencers, analog desks do not have a hard clipping threshold—levels can go above 0 dB. When planning levels on analog, engineers sometimes look at the VU meters—a far better indication of loudness than the peak meters on a software mixer. One might mix a dance track in Logic and start by setting the kick to –6 dB. Then, when another instrument is introduced, it might need some extra gain, and can even end up at the fader's dead-end. Employing a VU meter during the early stages of level planning can be beneficial in digital all the same.

The extremes-inward experiment

Sometimes level decisions are hard to make and we find it difficult to ascertain how loud a specific instrument should be. The extremes-inward experiment can help in these situations. Accompanied by Figure 13.7, the process is as follows:

- Take the fader all the way down.
- Bring it up gradually until the level seems reasonable.
- Mark the fader position.
- Take the fader all the way up (or to a point where the instrument is clearly too loud).
- Bring it down gradually until the level seems reasonable.
- Mark the fader position.
- You should now have two marks that set the limits of a level window. Now set the instrument level within this window based on the importance of the instrument.

If the result of the experiment is that the window is too wide, say more than 6 dB, it suggests that some compression or equalization might be beneficial.
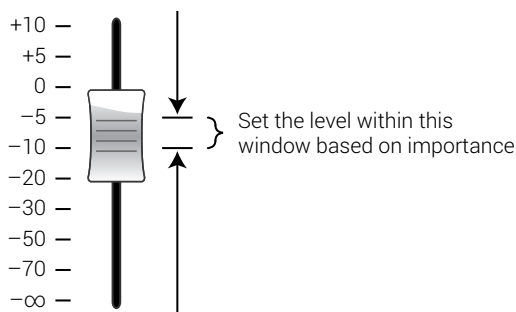
Figure 13.7 An illustration of the extremes-inward experiment.

14 Panning

How stereo works

On December 14, 1931, Alan Dower Blumlein, then a researcher and engineer at EMI, applied for patent number 394,325, called "Improvements in and relating to Sound-transmission, Sound-recording and Sound-reproducing System." This 22-page application outlined his ideas and vision to create a better sound reproduction system than the monophonic one used in those days. What Blumlein described as "binaural sound" is what we refer to today as stereo. His original concept was so far ahead of its time that many people could neither understand it nor realize its potential. The first stereo record was published in 1958—sixteen years after Blumlein's mysterious death and six years after EMI's patent rights had expired.

The term *binaural* denotes having or involving two ears. Our brain uses differences between the sound arriving at the left and the right ears to determine the localization of the sound source. The differences fit three criteria: amplitude, time (phase), and frequency. For example, if a sound is emitted from a trumpet placed to our right (Figure 14.1), the sound arriving at our right ear will be louder than that arriving at our left ear. This is due to the fact that the head absorbs and reflects some of the sound energy traveling to the left ear. As sound takes time to travel (approximately one foot per millisecond), it will also reach the left ear slightly later. Finally, since high frequencies are not very good at diffracting, they will not bend around our head like low frequencies, and less will reach the left ear.
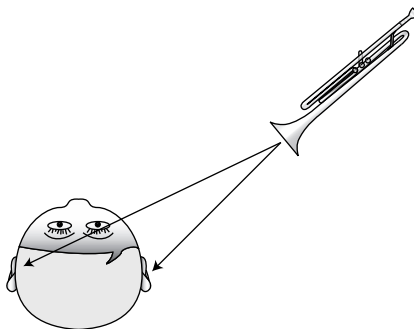
Figure 14.1 Sound produced from the right of the listener will have amplitude, time, and frequency differences between the left and right ears.

In order to simulate the binaural differences that occur in nature, we can use a dummy head with microphones at its ears or use a computer to do the calculations instead (this is how 3D effects are implemented). Either way, if we play such binaural material through headphones, we achieve a good sense of localization, with the ability to have sounds appearing as if coming from all around us, including the sides, behind, and even up and down. Essentially, we fool our brain by providing two distinct sounds with the same differences that sounds in nature involve.

Track 14.1: Congas Circle
A demonstration of binaural processing. Using headphones, the congas should appear as if circling the head. Listening to this track through speakers would only make the congas shift between the right and left channels.

Plugin: Wave Arts Panorama
Percussion: Toontrack EZdrummer

Using two speakers instead of headphones does not achieve the same effect. In nature, the sound from a central source in front of the listener travels an equal path, thus reaching both ears at the same time, at the same amplitude and with the same frequency content. But a two-channel stereo setup does not have a center speaker, so in order to create a central image both speakers emit the same sound. Having no real center speaker, this central image is known as *phantom center*. If we could ensure that the sound from each speaker only arrived at the nearest ear, as in Figure 14.2a, the simulation would be perfect. But the sound from each speaker also arrives to the far ear and does so slightly later, while being quieter and having less high frequencies (Figure 14.2b). This late arrival confuses our brains and results in a slight smearing of the perceived sound image. This smearing effect can be demonstrated in every 5.1 studio by comparing the sound emitted solely from the center speaker (real center) and that emitted at equal levels from both the left and right speakers (phantom center). How unfocused the latter might be can be quite surprising, especially in poorly tuned surround studios.

The best stereo perception is experienced when a listener sits on the central plane between a correctly configured pair of speakers. Sitting in such a position is requisite

(a)   (b)
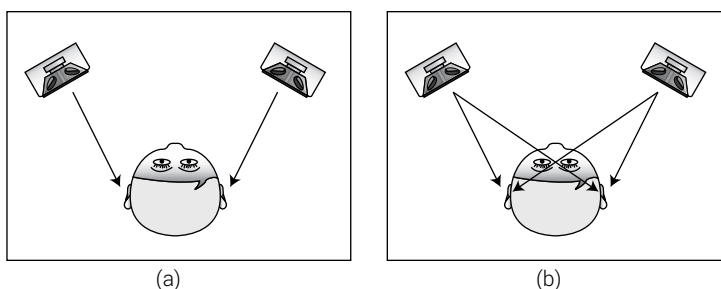
Figure 14.2 (a) A great stereo perception would be achieved if the sound from each channel reached its respective ear only. (b) In practice, sound from both speakers arrives at both ears, yielding a less accurate stereo image.

for any mixing engineer, but most people listen to music in far less ideal locations. Although some of the stereo information is retained from the sweet spot, much of it is lost, especially when moving away from the central plane.

Pan controls

The pan-pot

The first studio to install a stereo system was Abbey Road in London. Many other studios followed suit and soon enough all consoles had to accommodate this new feature. At first, consoles offered a three-state switch that could pan a mono signal either left, center, or right. This type of switch is evident in some old mixes where drums are panned to one extreme and bass to the other, and in all sorts of other panning strategies that are today only employed as a creative effect.

While working on *Fantasia* in the late 1930s, Walt Disney's engineers wanted to create the effect of sound moving between the two sides of the screen. Following on from the work of Blumlein and Fletcher (the same Fletcher from the Fletcher–Munson curves, who also researched multichannel sound at Bell Labs), they found that as the level of one speaker continuously drops in relation to the other speaker, the image gradually shifts toward the louder speaker. Studio engineers could use this knowledge. They could feed the same signal into two channel strips, pan each to a different extreme, then attenuate one channel. Depending on the amount of attenuation, engineers could pan instruments across the whole stereo panorama (Figure 14.3). The problem was that, with the small consoles used back then, this was a huge waste of a channel. Thus, the *panoramic potentiometer* was invented, otherwise simply known as the *pan-pot*.
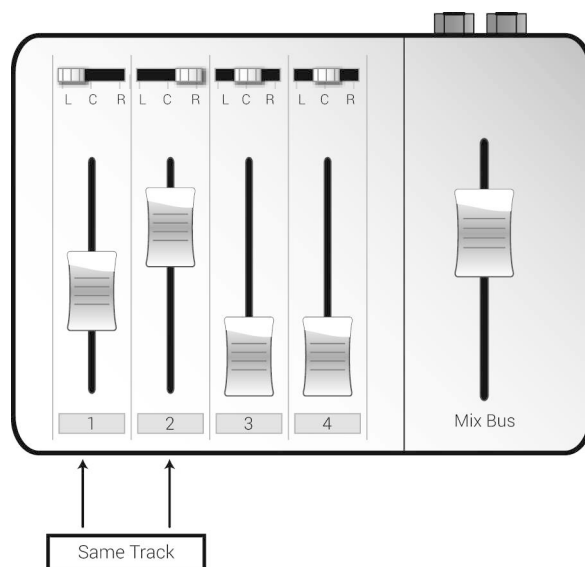
Figure 14.3 Manual panning. The same track is sent to two channels, each panned to a different mix extreme. When the signal that travels to one speaker is attenuated, the resultant image shifts to the other speaker.
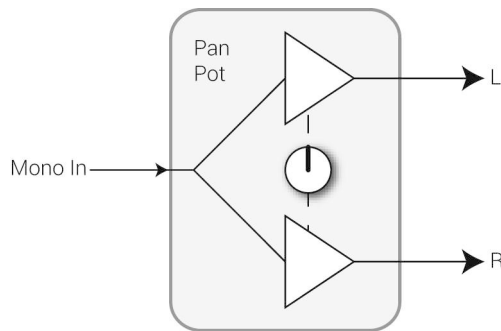
Figure 14.4 Inside a pan-pot. The signal is split and each copy is sent to a gain control. The pot itself only controls the amount of attenuation applied on one channel or the other.

A pan-pot does something very similar to the setup in Figure 14.3. It splits a mono signal to left and right channels, and attenuates one of the channels by a certain amount. Effectively, this alters the relative level between the two speakers. Figure 14.4 illustrates what happens inside a pan-pot. Panning something to one speaker only requires complete attenuation of the opposite channel. Center panning is achieved by sending the signal at equal levels to both speakers.

The pan clock

Often hours are used to describe the position of a pan-pot. Although no convention exists, often the lingo involves hours ranging from 7:00 to 17:00, with 7:00 being the left extreme (also known as *hard-left*), 12:00 being center, and 17:00 being the right extreme (*hard-right*). Some use a 12-hour clock, representing 17:00 as 5:00, but for clarity the 24-hour format will be used in this text. Also, some call the center *dead center*; in this text, the term *center* signifies 12:00, and any panning position around it will be referred to as "around the center" (or "slightly off-center").
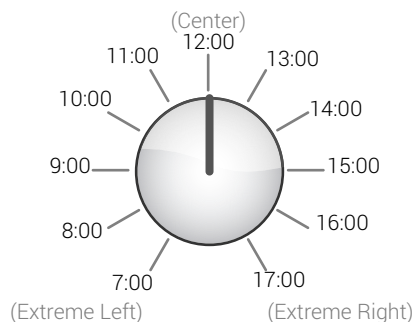
Figure 14.5 The pan clock.

Different software sequencers use different numeric scales to represent pan positions. Pro Tools and Cubase use a 201-step scale, while Digital Performer and Logic use the MIDI-based 128-step scale. Where hours might appear between the two speakers, and their equivalent values on the numeric scales, is shown in Figure 14.6.
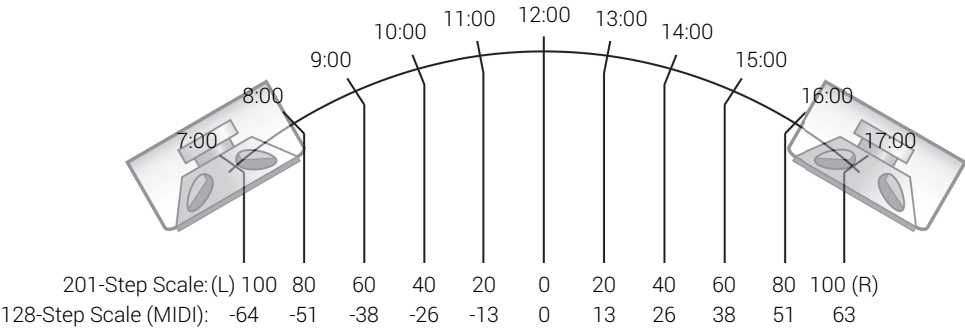
Figure 14.6 The rough position of the hours on the stereo panorama, and the hour equivalent numeric values on the 201-step and 128-step scales.

Track 14.2: Pan Clock
This track involves 11 pink noise bursts, each panned to a different hour starting from hard-left (7:00).

Panning laws

Not all pan-pots behave the same. Depending on the pan law, the signal level might rise, remain constant, or drop as we twiddle the pot from the extremes toward the center. Just to complicate things, the panning behavior is different in stereo and mono. A console usually offers one pan law for all of its pan-pots, although some in-line desks have two different pan laws—one for the channel path and another for the monitor path. Some applications let the user choose the pan law used in the software mixer (Figure 14.7).
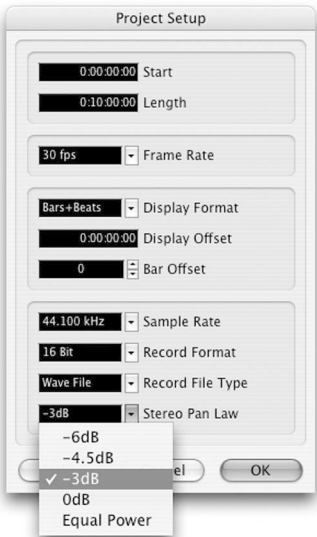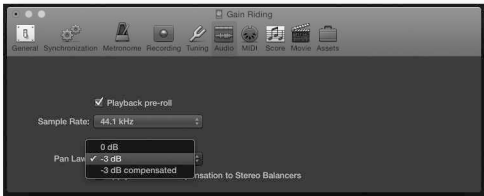
Figure 14.7 The pan law settings in (a) Cubase and (b) Logic.

In these applications, the choice of pan law is done per project and not globally for the system (to keep backward compatibility). Changing the pan law halfway through a mix would alter the relative levels between instruments, so a choice must be made before mixing commences.

There are two main principles to be aware of in panning law. First, acoustics has it that if two speakers emit the same signal at the same level, a listener on the central plane will perceive a 3 dB boost compared with what each speaker produces. For example, if two speakers emit 80 dBSPL each, the perceived loudness will be 83 dBSPL. The second principle is concerned with how mono is achieved. When two channels are summed to mono, half the level of each is sent to both. For example, with a signal panned hard-left, the left speaker might produce 80 dBSPL and the right speaker 0 dBSPL. Half the level is approximately –6 dB, so when summed to mono each speaker will produce 74 dBSPL. Based on the first principle, the perceived loudness in this case will be 77 dBSPL. We can see from this example that if a signal panned to one extreme is played in mono, it drops by 3 dB. For centrally panned signals, mono summation makes no difference—half the level of each channel sent to both channels results in exactly the same level on both.

Thus, the level of centrally panned signals remains consistent whether played in stereo or mono. The fact that when summed to mono the center mix-image remains unaffected and the extremes drop by 3 dB is not such a bad thing if we consider that most of the important instruments are panned center.

> *When summed to mono, the center image remains at the same level while the extremes drop by 3 dB.*

The four common pan laws, –0, –3, –4.5, and –6 dB, are presented below for both stereo and mono. Before delving into each pan law, it is worth understanding, as mixing engineers, what it is we really care about. Depending on the pan law, mix levels might change as we pan. We generally do not want this to happen while we mix, as each pan move might call for subsequent fader adjustments. Since we mix in stereo, our only true concern is for levels to stay uniform as we pan across in stereo. Once the mix is bounced, the pan law is locked—the level balance printed on the master is determined and final.

> *How different laws behave in mono is of great importance in radio drama—signals panned from one extreme to another (for instance, a character crossing the stage) will be perceived differently on stereo and mono receivers.*

Figure 14.8 illustrates how the perceived level varies for different pan positions with the **0 dB pan law**. Pan-pots of this kind do not drop the level of centrally panned signals. A signal panned hard-left might cause an 80 dBSPL radiation from the left speaker. If the signal is then panned center, both speakers will produce 80 dBSPL, which results in a 3 dB increase in perceived loudness. An instrument will rise in level as we pan it from either extreme toward the center (in other words, the instrument
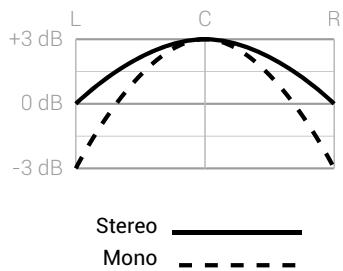
Figure 14.8 The perceived loudness of the 0 dB pan law.
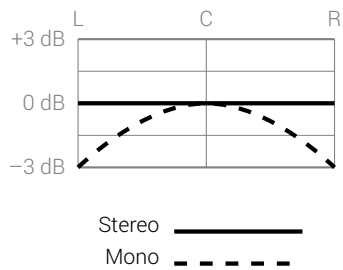
Figure 14.9 The perceived loudness of the –3 dB pan law.

level will drop as we pan from the center outward). While panning from left to right in stereo, an instrument will have a 3 dB center boost. In mono, which is not of much interest to us, the standard 3 dB extremes drop happens, and while panning from left to right an instrument will experience a 6 dB center boost.

The **–3 dB pan law** (Figure 14.9) compensates for the 3 dB center boost that occurs with the 0 dB law. Centrally panned signals on pots of this kind are sent to each speaker at –3 dB compared with the level of a signal panned to either extreme. Eighty dBSPL from one speaker for a signal panned hard-left will become 77 dBSPL from each speaker when the signal is panned center. This is a total perceived loudness of 80 dBSPL. When panning from left to right in stereo, signals remain at the same level. This is highly desirable while mixing, as pan adjustments rarely bring about the need for subsequent fader adjustments.

With the –3 dB pan law, signals panned from left to right have a uniform level in stereo, but in mono the perceived level has a 3 dB center boost. For mono-critical applications, the –6 dB law in Figure 14.10 is used. It provides a uniform level in mono, but a 3 dB center dip in stereo. Similarly to the 0 dB law, the –6 dB law is not recommended for stereo mixing, as fader adjustments might be needed after panning. Another pan law, the –4.5 dB, is a compromise between the –3 and –6 dB laws. In stereo, there is a 1.5 dB center dip, while in mono there is a 1.5 dB center boost (Figure 14.11). While neither in stereo nor in mono, the signal level is consistent as it is panned across; the maximum error is 1.5 dB and not 3 dB, as with the –6 and –3 dB laws.
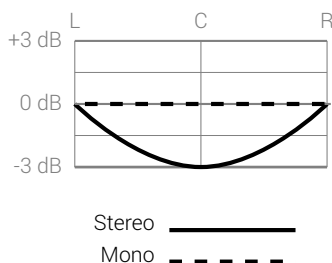
Figure 14.10 The perceived loudness of the –6 dB pan law.
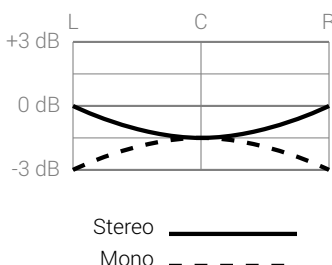
Figure 14.11 The perceived loudness of the –4.5 dB pan law.

Other variables might affect the choice of pan law. The first principle we discussed tells us that the power summation of two speakers produces a 3 dB increase on the central plane. In practice, the resultant loudness increase in less reverberant rooms might be 6 dB, but for low frequencies only. Having a low-frequency boost of 6 dB and a high-frequency boost of only 3 dB makes the –4.5 dB law a reasonable choice. Another issue is based on the assumption that most masking occurs in the central area, which means that panning instruments to the sides can result in an apparent loudness increase. The –2.5 dB pan law has a subtle +0.5 dB center boost to compensate for such a phenomenon. (For most of its history, Pro Tools' pan-pots used the –2.5 dB law; only recently other pan laws became available.) When a pan law choice is available, experiments can help us determine which pan law brings about uniform levels as sounds are panned across. The –3 dB law is very likely to be what we are after.

The –3 dB pan law is generally the best option for stereo mixing.

The following tracks demonstrate the different pan laws. Pink noise sweeps across the stereo panorama, and for each pan law two versions are produced (stereo and mono) and then bounced. Perhaps the most important thing to observe is how the level varies on the stereo versions as the noise is swept between the center and the extremes.

Track 14.3: Pan Law 0 dB Stereo

Track 14.4: Pan Law 0 dB Mono

Track 14.5: Pan Law –3 dB Stereo

Track 14.6: Pan Law –3 dB Mono

Track 14.7: Pan Law –4–5 dB Stereo

Track 14.8: Pan Law –4–5 dB Mono

Track 14.9: Pan Law –6 dB Stereo

Track 14.10: Pan Law –6 dB Mono

The balance pot

Unlike the pan-pot, the input to a balance pot is stereo, as can be seen in Figure 14.12. The two input channels each pass through a separate gain stage before being routed to the stereo output. The pot position determines the amount of attenuation applied on each channel. It is important to note that a balance pot never crossfeeds the input signal from one channel to the output of another. This results in a serious mixing limitation—the image of the input stereo signal is tied to at least one extreme. This is illustrated in Figure 14.14. We can narrow the stereo width of the input signal, but we cannot set it free from ending at one extreme.
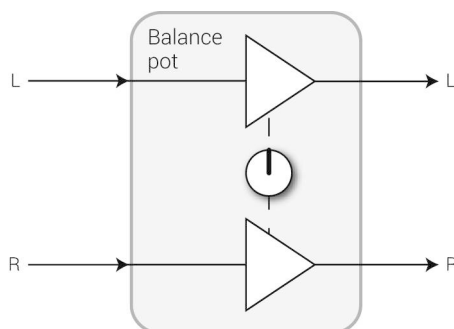
Figure 14.12 Inside a balance pot. Both the left and right inputs pass through separate gain stages. But the input from one channel never blends into the output of the other.
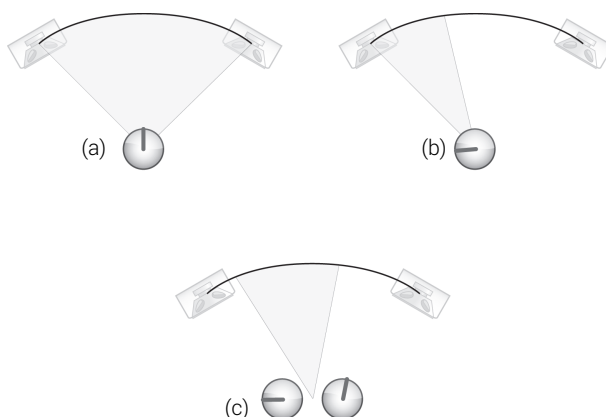
Figure 14.13 Both (a) and (b) show the resultant image of a specific balance position. As can be seen in (b), when the balance pot is turned left, the span of the resultant stereo image narrows but remains in the extreme of the stereo panorama. (c) If we need to pan anything so the span of the stereo image does not remain in the extremes, we need two pan-pots instead of a single balance pot.

(a)   (b)
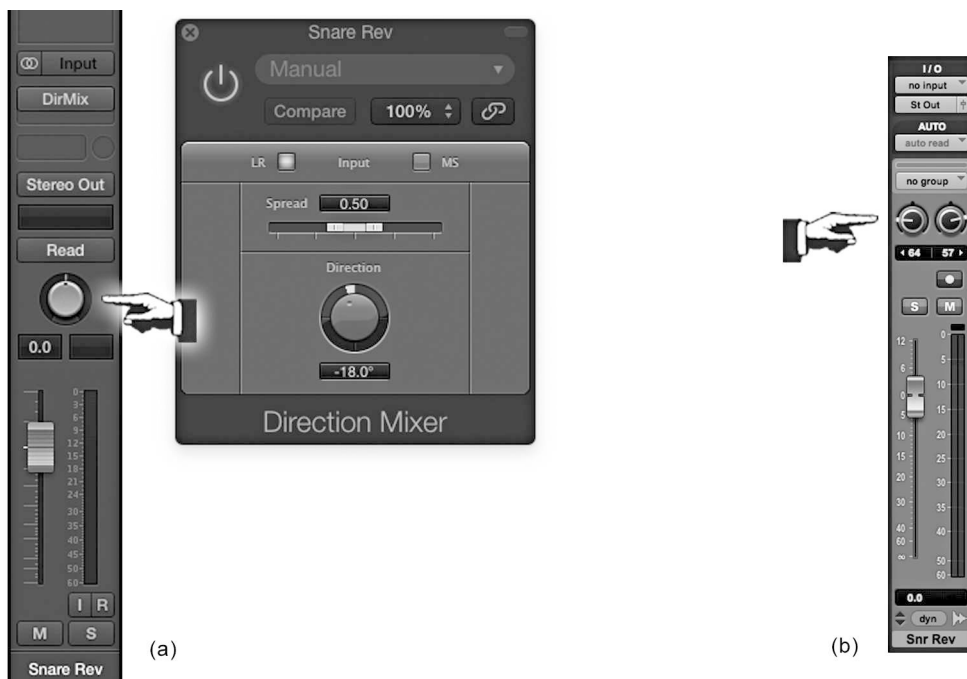
Figure 14.14 (a) Logic only offers balance control for stereo tracks. More control over stereo image can be achieved with the help of the direction mixer plugin. (b) Pro Tools offers dual-mono pan control for stereo tracks, which provides full functionality for any stereo imaging task.

In order to narrow the width of a stereo signal and place it freely across the stereo panorama, two pan-pots are required (Figure 14.13c). Unfortunately, not all software mixers provide a dual-mono control for stereo tracks. Applications such as Logic only provide a balance control, and users need to load a plugin in order to place stereo signals freely across the stereo panorama (Figure 14.14).

Track 14.11: Balance Stereo Delay
This track involves a synth sent to a stereo delay.

Track 14.12: Balance Hard-left
The balance pot on the delay return is panned hard-left. As a result, the right channel of the delay is lost.

Track 14.13: Balance Hard-right
The balance pot on the delay return is panned hard-right. As a result, the left channel of the delay is lost.

Track 14.14: Balance No Can Do
Instead of using a balance pot on the delay return, a dual-pan control is employed. The left and right channels are panned to 10:00 and 14:00, respectively, narrowing the echoes toward the center. No balance control can achieve this, alas.

Plugin: PSP 84

Types of track

A single-instrument recording can involve more than one track. Raw tracks can be classified like so:

- Mono
- Stereo:
    - Acoustic:
        - Coincident pair (XY)
        - Spaced pair (AB)
        - Near-coincident pair
    - Electronic
- Multiple mono

Mono tracks

A single instrument, a single take, a single track. A mono track can be panned anywhere across the stereo panorama. The main problem with a dry mono track is that it provides no spatial cue. Only in an anechoic chamber will signals arrive at our ears without some room response—without some space being noted. When a dry mono track is placed untreated, it can appear foreign to the depth and ambiance aspects of the mix, and can appear frontal. A reverb emulator, or other spatial effect, is commonly employed in order to blend mono tracks into the spatial field of the mix. Even a tiny amount of ambiance reverb can do the job. A real problem can arise when a mono track includes some room or artificial reverb. The monophonic reverb is unlikely to mix well with the stereo reverb of the mix ambiance. Sending the track, with its embedded reverb, to the ambiance reverb might not be wise—a reverb of a reverb seldom makes any spatial sense.

Mono tracks often benefit from the addition of some spatial effect that blends them into the ambiance of the mix.

Track 14.15: Drum Loop
The loop under discussion in the following two tracks.

Track 14.16: Drum Loop Dry
When the drum loop is mixed dry, it can appear foreign to the mix—it is focused and positioned at the very front of the sound stage.

Track 14.17: Drum Loop with Reverb
By adding some reverb, the drum loop blends into the mix better, and it is not as distinguishable as in the previous track.

Plugin: Audio Ease Altiverb

Stereo pairs

A stereo pair track denotes a pair of mono tracks; each represents one of two microphones arranged in an established stereo miking technique. The two microphones

are positioned in a way that will cause amplitude and time differences between the signals captured by each microphone. When the microphones are panned to the opposite extremes, the differences between them result in spatial reproduction. Stereo miking techniques are divided into three families: *coincident pair*, *spaced pair*, and *near-coincident pair*. While mostly a recording topic, each technique affects how we mix the given tracks.

The **coincident pair** technique, also known as XY, is based on two microphones angled apart with their diaphragms nearly touching. All but sounds arriving from the center will present a different amplitude on each microphone. Because of the close proximity of the diaphragms, phase differences are disregarded. Due to the phase similarity between the two microphones, this family is known to provide the **best mono-compatibility**. The two tracks can be summed internally without causing noticeable comb filtering, and thus the two tracks can be panned inward from the extremes. This is often done in order to narrow the image width of an instrument, like with drum overheads.

> *Coincident pair gives the best mono-compatibility, which allows us to narrow the image of the stereo source during mixdown.*

The coincident pair family includes stereo miking techniques such as Blumlein and MS (both devised by Alan Blumlein) or Power-thrust (hyper-cardioids at 110°). Of all these techniques, MS (mid-side) is most worthy of discussion, as it requires a special decoding matrix. The technique is based on a directional microphone facing the sound source (M) and a bidirectional polar pattern picking up the sides (S). Only two tracks are recorded, but on a desk these are decoded using three channels. The M is routed to one channel panned center. The S is routed to two oppositely panned channels, one of which is phase-inverted (Figure 14.15). The relative level of S determines the amount of reverberation and the stereo spread of the source in the mix.

More on MS in Chapter 27.

The **spaced pair** technique is also known as AB (often incorrectly titled XY). The arrangement of a spaced pair involves two microphones spaced a few feet apart; as a result of this distance, phase differences do occur. The microphones are usually omnidirectional so amplitude differences can be disregarded. Due to the phase differences between the two microphones, this technique is characterized by increased spaciousness and vague imaging. However, the same phase differences mean that this technique is **not monocompatible**. Any inward panning of the two microphones might result in noticeable comb filtering, sometimes forcing the two microphone tracks to be panned to the opposite extremes.

The coincident pair family is based on amplitude differences, while the spaced pair family is based on phase differences. The **near-coincident** family is a marriage of both—two microphones are angled and spaced, but only a few inches apart. The phase differences of a near-coincident recording do not present the same mono-compatibility
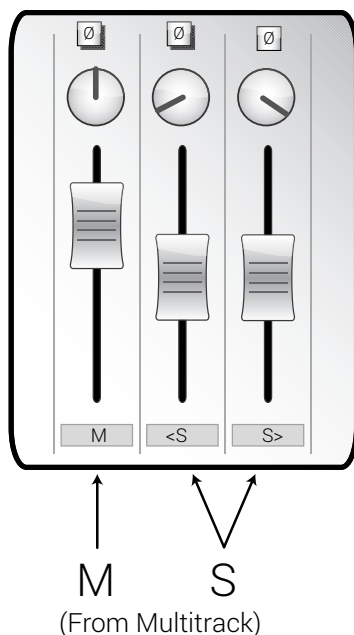
Figure 14.15 The MS technique. The M is routed to one channel panned
center. The S is routed to two oppositely panned channels, one of which is
phase inverted.

as a coincident pair, but the differences aren't as profound as with a spaced pair.
Panning the two microphones to anywhere but the extremes can result in some degree
of comb filtering, which may or may not be noticeable.

In addition to the acoustic stereo pair, we can also get a stereo recording of a synthesizer
or other electronic instrument. Whether or not the image of these can be narrowed can
only be determined through experiment. But it is worth knowing that, despite having a
stereo output, some of these devices are monophonic inside. The stereo size these
devices produce might be achieved by having the pure monophonic output on the left
channel and a processed version of it on the right channel. This is known as *pseudo-
stereo*. When the stereo effect fails to blend into the mix, we can always abandon the
processed channel and apply our own effect on the pure monophonic signal.

Multiple mono tracks

Multiple mono tracks are those that represent the same take of the same instrument,
but do not involve an established stereo miking technique. Often each track captures a
different timbre component of the sound source, so during mixdown we can shape the
instrument sound using faders rather than processors. Some of these complementary
tracks would be panned to different positions, most likely to the two extremes. An
example would be two acoustic guitar tracks—soundhole and neck.

More often, though, these complementary tracks are panned to the same position, a
practice that requires them to be in phase with one another. Examples include snare-top
and snare-bottom, bass-mic and bass-direct, or three microphones used to capture the
full timbre of a double bass. However, identical panning positions are not obligatory.

When appropriate, the tracks can be panned to different positions, which will widen and blur the stereo image of the instrument. For instance, we can achieve such a blurry, wide snare image by panning the snare-top and snare-bottom differently. The effect can be made subtler if one of the tracks is attenuated. This effect can also be achieved artificially by panning a duplicate of a mono track to a different position, while generously altering its tonal character.

Track 14.18: Snare Top Bot Same
In this track, the snare-top and snare-bottom microphones are panned to the same position. The snare image is not perfectly focused due to the overheads and room-mic, but its stereo position can easily be pointed out.

Track 14.19: Snare Top Bot Mirrored
The snare-top is panned to the same position as in the previous track, but the snare-bottom is panned to the opposite side. This results in a less focused snare image, and an exact position is harder to identify.

Drums: Toontrack EZdrummer

Just as not all the songs recorded for an album are actually used, not all multiple mono tracks must be used. Sometimes using both bass-mic and bass-direct calls for treating each individually and then both collectively—this might involve two equalizers, three compressors, and a control group. Using one track only might not only be easier, but can also fall perfectly into the requirements of the mix. In the case of guitar-hole and guitarneck tracks, sometimes excluding one track and adding reverb to the other works better than mixing both tracks.

Combinations

With some recordings, such as drums or orchestral, we get a blend of different types of tracks; for instance, stereo overheads, multiple mono kick, and mono toms. In such cases, the stereo pair provides a reference image to the location of its constitutes. The respective mono tracks are often panned to the same position. For example, if the snare appears to come from 11:00 on the overheads, the snare track might be panned to that position. Again, this is not a hard-and-fast rule—different panning will result in wider, less focused images, which is sometimes sought (Figure 14.16).
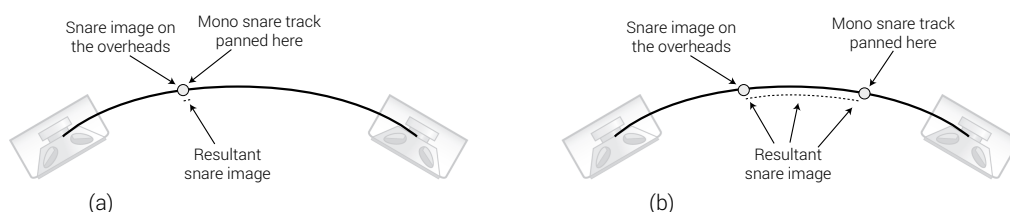
Figure 14.16 The stereo image of a snare, depending on how its mono track is panned compared to the overheads' image. (a) The snare is panned to the same position as it appears in the overheads' stereo image, resulting in an overall focused image for the snare. (b) A snare panned far away from its overhead position will have a wider, unfocused overall stereo image, which is sometimes desired.

Panning techniques

An initial panning plan can be drawn just by looking at the track sheet. The nature of each track will usually hint at its rough panning position in the mix. The following is an example of an initial panning plan for a simple rock production:

- **Overheads**—70 percent wide around the center. Audience view.
- **Kick**—center.
- **Snare**—to the same location as it appears on the overheads.
- **Snare reverb**—40 percent wide, around the snare position.
- **Tom 1**—14:00.
- **Tom 2**—13:00.
- **Tom 3**—10:00.
- **Bass**—center.
- **Vocals**—center or maybe slightly off-center to the left.
- **Power guitar I**—hard-left (+Haas trick).
- **Power guitar II**—hard-right (+Haas trick).
- **Flutes—**halfway toward the extremes.
- **Ambiance reverb**—100 percent stereo width.

Plans are never successful until executed. Clearly, the above arrangement will not be suitable for all productions; it is just a possible starting point. A calculated approach, however, can backfire—different panning schemes can project very different feels, sizes, and power. In complex arrangements, experiments may be a necessity, and part of our role is to explore, discover, and reflect on different panning strategies.

Panning schemes can have a profound effect on the mix.

As the mix progresses and elements are added, we revisit the pans. Sometimes even small alterations result in big improvements—just shut your eyes and inch it until it latches into place. As already mentioned, panning alterations are sometimes part of the very last refinements of a mix.

Panning and masking

We already know that internal summing is less forgiving compared with acoustic summing. When two instruments are panned center, their internal summing accents any masking interaction between them. However, if each is panned to a different extreme, they are summed acoustically, which makes masking interactions less harmful. As a result, panning instruments toward the extremes also increases their definition. On the very same basis, masking is more dominant when two fighting instruments are panned to the same side of the panorama; mirroring one of them (e.g., from 15:00 to 9:00) can improve the definition of both.

The sound stage

Essentially, the stereo panorama pertain the horizontal plane of the imaginary sound stage. One of the main roles of pan-pots is to determine where instruments are positioned on that stage. Just like on a real stage, listeners expect to find the most important instrument in the center of the mix—be it the vocalist or the trumpet

player. Also, if we imagine a five-piece rock band on a big festival stage, it is unlikely that the two guitarists will stand on the stage edges; it is more likely that each will be somewhere halfway along each side. Many more examples can be given, but contemporary music is not always mixed to the full realistic extent of a live performance—if the guitars sound better when panned each to a different extreme, one would be a fool not to pan them that way. The visual sound stage is a guide often more suitable for natural mixes.

Low-frequency instruments, notably the bass and kick, are usually panned center. One reason is that low-frequency reproduction demands much more power than the reproduction of high frequencies. Panning low frequencies off-center will cause uneven power consumption between the two channels, which can have many downgrading consequences on the combined stereo response. The way vinyls are cut is another reason. What's more, panning low-frequency instruments to the sides is somewhat pointless since low frequencies give very little directional cue (hence subwoofers can be positioned off the central plane). Yet it is worth remembering that both kicks and basses have important content in their mids as well.

Having both the most important and the low-frequency instruments panned center makes it inevitably the busiest area of most mixes, and where most masking happens. No panning police will arrest you for not panning the most important instruments dead center. Lead vocals, for example, can be panned slightly off-center; sometimes even the bass and kick are panned that way. Most listeners do not listen on the central plane anyway, and even if they do they are unlikely to recognize these subtle panning shifts.

*The center is usually the busiest area in the mix. Instruments that would
usually be panned center can be panned slightly off-center.*

Track 14.20: Vocal Center
The vocal on this track is panned center (12:00).

Track 14.21: Vocal 1230
The slightly off-center (12:30) panning of the vocal in this track can easily go unnoticed, especially if the track is not compared to the previous one.

Track 14.22: Vocal 1300
The appreciation of the vocal panning in this track is dependent on the width of the monitoring setup. While mixing engineers with a wide stereo setup could find this off-center panning disturbing, normal listeners could easily fail to notice the off-center vocals.

Track 14.23: Vocal 1400
With the vocal panned at 14:00, it is hard to overlook the off-center image of the vocal. Both engineers using a narrow stereo setup and normal listeners could find this panning scheme disturbing.

Another important area of the mix is the extremes. We have said already that overpanning to the extremes can result in mixes that have busy extremes and center but nothing in between; these are known as W-mixes. W-mixes are the outcome of the

tendency to pan every stereo signal and every pair of mono instruments to the ex-
tremes. While it could be beneficial to pan overheads, stereo effects, or a pair of mono
guitars that way, most mixes will benefit from a **stereo spread balance**, where mix
elements are panned to fill the sound stage.

Not all stereo elements and mono pairs should be panned to the extremes.

It is a common practice in surround film-mixing that the voice of an actor standing
behind the camera is still panned to the front speakers. Had it come from the rear,
viewers would notice the existence of the rear speakers, which might distract them
from the visual experience. As a general rule, cinema sound should accompany the
visual but never draw too much attention on its own—you should feel as if part of the
scene, but should not notice how this is done technically. The same principal applies
with music mixing; the listener's realization that two distinct speakers are involved in
the production of a sonic experience can distract from the spatial illusion. Highly cre-
ative panning techniques should be used sparingly and only when appropriate.

Another fact worth remembering is that the center of the mix is the least-focused area
(due to the late arrival of sound from both speakers to the respective far ear), while the
extremes provide the most-focused area (as only one speaker generates the sound in true
mono). Consequently, as instruments are panned toward the extremes, they become
more focused. Combined with the curved nature of the mix front line, these instruments
can also appear closer. This implies that the more an instrument is panned toward the
extremes, the more it might benefit from the addition of a spatial effect—especially in the
case of mono tracks.

The extremes = more focus, more definition, closer image.

*When comparing the following tracks, the drum loop should appear closer and more distinct
and focused when panned to the right extreme:*

Track 14.24: Drum Loop Panned Center

Track 14.25: Drum Loop Panned Hard-right

As discussed in Chapter 7, one of the main mixing objectives when it comes to the
stereo domain is balance. Both **level and frequency** balance are our main concerns.
The off-center panning of individual instruments is guided by this objective, and rough
symmetry is often sought. One guitar to this side, another guitar to the other; a shaker
to this side, a tambourine to the other; first vocal harmony at 11:00, the second at
13:00. A stereo balance relies on both instrumentation and recordings. "What instru-
ment can we record that complements a shaker?" or "Shall we double-track the
shaker?" are questions that should be asked before the mixdown stage. When an off-
center track creates imbalance that no other track can repair, mixing engineers often
utilize stereo effects, for example by sending a delayed echo to the opposite of a dry
instrument (while perhaps altering its tonality).

We can also associate stereo balance with **performance**. Say 16th notes are played on the
hi-hats and there are two guitars—one plays a 16th-note arpeggio, the other half-notes.

Panning the hi-hats and the arpeggio guitar to opposite sides will create rhythmical balance. But panning the hi-hats and the arpeggio guitar to the same side and the half-notes guitar to the other will result in one extreme being rhythmically faster than the other. Subsequently, this will also affect the stereo balance between the two sides of the stereo panorama.

While balance is a general stereo objective, a perfectly stereo-balanced mix is essentially a mono mix. There is always a margin for some degree of disparity. If we take a drum kit, for example, with its hi-hats and floor tom on opposite sides, imbalance is inevitable. But if a tambourine also forms part of the arrangement, it would make more sense to pan it opposite the hi-hats. As per our axiom that "percussives weigh less," sustained performance is more of a factor when it comes to stereo balance.

Effect panning

Stereo effects can be panned in various ways. The panning choice is mostly based on two aspects: width and center position. An ambiance reverb is one type of stereo effect that can benefit from full-width panning, especially when natural results and a wide sound stage are sought. But reverbs and delay lines used for creative purposes can benefit from a narrower image and off-center position. A classic example is the snare reverb, which is often different from the reverbs used for ambiance (or from the ambiance on the recording). When panned to the extremes, the wide snare reverb might compete with the ambiance and interfere with the spatial effect the latter provides. Combined, it might sound like the mix took place in a hall, while the snare was in a bathroom. Narrowing the reverb width will give a clear cue that it is not part of the general ambiance, but only an effect linked to the snare. This also promotes separation.

Matching the center of the snare and its reverb will emphasize this relationship. A mix with a snare panned to 11:00 might benefit from the snare reverb being panned around it, say at 10:00 to 12:00. Narrowing the width of a stereo effect can also be beneficial with stereo delays. The combination of centrally panned vocals with hard-panned echoes makes the echoes noticeable, spatially random, and distant from their parent. To make the effect more subtle and focused, we can pan the echoes inward toward the location of the vocals.

Tracks 22.38–22.40 in Chapter 22 demonstrate various stereo delay-panning techniques. Tracks 24.86–24.93 in Chapter 24 demonstrate various reverb-panning techniques.

Mono effects can benefit from the opposite way of thinking. The closer the effect is to its parent instrument, the less noticeable it will be, and we miss an opportunity to balance the stereo panorama. An example could be a mono delay applied on hi-hats. If the delay is panned to the same position as the hi-hats, the two are likely to collide. Panning the delay to the opposite side will increase the definition of both, while enriching the stereo panorama. One of the world's most famous delay lines is said to have been created by accident when Giorgio Moroder panned a bassline to one speaker and its delay to the other. This effect, which originally appeared on Donna Summer's "I Feel Love," has been popular ever since, and can be heard on countless dance mixes.

Track 14.26: Moroder Delay

As with "I Feel Love," the bassline (8th-notes sequence) and the hats are panned hard-left. The 16th-note delays of both the bassline and hats are panned hard-right.

Track 14.27: Moroder No Delay

Same as the previous track but without the delay.

Track 14.28: Moroder Delay Only

In this track, the source bassline and the hats are muted, but not their delay. The delay is evidently a 16th-note behind the main beat.

Track 14.29: Moroder All Center

This is the result of panning all tracks, including the delay, to the same position. The effect is distinct from that on Track 14.26.

Plugin: PSP 84

Beyond pan-pots

Autopanners

An autopanner pans audio back and forth between the left and right sides of the stereo panorama in a cyclical way. The operation of an autopanner might involve these parameters (Figure 14.17):

- **Rate**—when defined in Hz, determines the number of cycles per second; each cycle consisting of the audio being panned from the center to one extreme, to the other, and back to the center. Alternatively, the duration of a cycle might be defined in a tempo-related unit, such as a bar or half-note. With a slow rate (roughly below 1 Hz), the left and right movement is clearly recognizable. This recognition is lost with higher rates, where the effect resembles a Leslie cabinet.

Figure 14.17 The MOTU *AutoPan* plugin.

- **Depth**—defines in percentage how far toward the extreme the signal will be panned. With 100 percent, the varying panning position will reach the extremes; 50 percent halfway to the extremes; and so forth. The higher the depth setting, the more apparent the effect.
- **Waveform**—defines the shape of the pan modulation. A sine wave will vary the position between left and right in a sinusoidal way. A square wave will make the audio jump between the two sides.
- **Center**—defines the center position of the modulation. With the center set halfway to the left extreme, it is possible for the audio to only drift back and forth between the center and the left extreme.

Track 14.30: Beat Source
The source track used for the following tracks.

Track 14.31: Beat Autopan Tri Depth 100
Triangle waveform modulation, 100 percent depth, rate is tempo-synced to one bar (center, as with all other samples here, is set to 50 percent, which is 12:00).

Track 14.32: Beat Autopan Sine Depth 100
Sine waveform, 100 percent depth, rate is tempo-synced to one bar.

Track 14.33: Beat Autopan Sine Depth 50
Sine waveform, 50 percent depth, rate is tempo-synced to one bar.

Track 14.34: Beat Autopan Sine 5 Hz
Sine waveform, 100 percent depth, rate is set to 5 Hz. The cyclic left and right movement becomes vague due to the fast rate.

Track 14.35: Beat Autopan Sine 10 Hz
Sine waveform, 100 percent depth, rate is set to 10 Hz. The fast rate creates a unique effect, but one that doesn't resemble autopanning. It is interesting to note the kick, which appears at different positions every time it hits.

Plugin: MOTU AutoPan

The traditional autopanner effect—a slow, obvious, left-and-right movement—has lost popularity over the years, perhaps due to being gimmicky and not playing any vital mixing role. Today, the tendency is to use the effect subtly, or occasionally in specific sections of the song (mostly for interest's sake). An interesting motion effect can be achieved when more than one track is autopanned simultaneously, and all the tracks complement one another to produce a relatively balanced stereo image; for example, having three guitars moving across the stereo pano-rama, with each starting at a different position and perhaps autopanned with a different rate.

Autopanners can be connected to other processors in order to create cyclic morphing between two effects. Instead of routing the autopanner outputs to the mix, each can be routed to a different bus, which feeds a different processor. For example, the left output might feed a reverb, while the right output might feed a delay. This will make one effect appear when the other disappears, and vice versa. We can also use only one output of the autopanner for an effect that we simply want to come and go (however, a slow tremolo might be easier to set up for this purpose).

15 Equalizers

In the early days of telephony, the engineers at Bell Labs faced a problem: high frequencies suffered loss over long cable runs, a fact that would have made the voice during long-distance calls dull and hard to understand. They set to design an electronic circuit that would make the sound on both ends of the line equal. The name later given to this circuit was *equalizer*.

The equalizers used in mixing today are not employed to make one sound equal to another, but to manipulate the frequency content of various mix elements. The frequency virtue of each individual instrument, and how the latter appears in the overall frequency spectrum of the mix, is a paramount aspect of mixing. Operating an equalizer is easy, but distinguishing frequencies and mastering their manipulation is perhaps the greatest challenge mixing has to offer. This is worth repeating:

> *Distinguishing frequencies and mastering their manipulation*
> *is perhaps the greatest challenge mixing has to offer.*

Equalization and filtering is one area where digital designs overpower their analog counterparts. Many unwanted artifacts in analog designs are easily rectified in the digital domain. Digital equalizers provide more variable controls and greater flexibility. For example, it is very common for a digital equalizer to offer variable slopes—a rare feature of analog designs. Digital filters can also have steeper or narrower responses. Theory has it that whatever can be done in the analog domain can be done in the digital one, and vice versa. However, to have the same features as some digital equalizers, analog designs would end up being very costly and prohibitively noisy. This by no means implies that digital equalizers sound better than analog ones, but equalization is one area in mixing that has experienced a radical upgrade with the introduction of digital equalizers, and, more specifically, software plugins.

Applications

Although equalizers are not the only tools that alter frequencies, they are the most common. In simple terms, equalizers change the tonality of signals. This simple ability is behind many, many applications.

Balanced frequency spectrum

This is an ever-so-important aspect of a mix. It was suggested in Chapter 7 that, despite a vague definition of what tonal balance is, a deformed frequency response is unlikely to go unnoticed. Overemphasized or lacking frequency ranges, whether wide or narrow, is one of the critical faults that equalizers are employed to correct. They also help us to narrow or widen the frequency "size" of instruments, or shift them lower or higher on the frequency spectrum.

Shaping the presentation and timbre of instruments

Equalizers give us comprehensive control over the tonal presentation of instruments and their timbre. We can make instruments appear thin or fat, big or small, clean or dirty, elegant or rude, sharp or rounded, and so forth. A great example of this is when mixing kicks—equalizers can change their timbre in the most radical ways. Whether kicks or any other instrument, shaping the timbre and tonal presentation of instruments is one of the more creative aspects of mixing.

Separation

The frequency ranges of various mix elements nearly always overlap, a sub-bass (low-frequency sine wave) being a rare exception. Most instruments end up in a masking war that can, in some mixes at least, get quite bloody. When two or more instruments are fighting for the same frequency range, we can find it hard to discern one instrument from the other, as well as fail to perceive important timbre components or certain notes. As long as instruments are mixed together, they inevitably mask one another; we can ease the conflict by cutting from instruments any dispensable or otherwise less-essential frequencies, or in some cases by boosting important frequencies.

It is important to mention here that we do not limit each instrument to a unique frequency range—such a practice would make each instrument and the overall mix sound horrific. We combat masking until we can separate one instrument from another, and all instruments are defined to our satisfaction. Depending on the nature and density of the arrangement, this may or may not involve serious equalization.

Definition

Definition is a subset of separation. No separation equals low definition. But we also associate definition with how recognizable instruments are (provided we want them recognizable) or how natural they sound (provided we want them to sound natural). For example, in a vocal and piano arrangement there might be great separation between the two; but if the piano sounds as if it is coming from the bottom of the ocean, it is fair to say that it is poorly defined. The same can be said for a bass guitar if the notes being played cannot be discerned.

Conveying feelings and mood

Our brain associates different frequencies with different emotions. Bright sounds tend to convey a livelier, happy message, while dark sounds might be associated with mystery or sadness. Because the human voice deepens with adulthood, some believe

that low-frequency attenuation of the voice gives a more youthful impression. Equalizers can be used to make vocals sweeter, a snare more aggressive, a trumpet mellower, a viola softer, etc.

Creative use

Equalizers are not employed solely for practical reasons. Creative use entails less natural or "correct" equalization that can put instruments in the limelight, give a retro feel, or create many fascinating effects.

Interest

Automating EQs is one way to add some interest to a mix. A momentary telephone effect on vocals, the rolling off of lows during a drop, and the brightening of the snare during the verse are just a few examples.

Depth enhancements

Low frequencies bend around obstacles and are not readily absorbed. High frequencies do the opposite. These are some of the reasons why we hear more low-frequency leakage coming from within venues. Our brain decodes dull sounds as if coming from farther away, which is why our perception of depth is stretched when underwater. We use this darker-equals-farther phenomenon to enhance, if not to perfect, the front/back localization of instruments in the mix.

Track 15.1: EQ and Depth
This track is produced with an LPF (low-pass filter) swept down and then up. Note how the fewer high frequencies there are, the further back the drums appear. (This is also the outcome of the level reduction caused by the filter.)

Plugin: Digidesign DigiRack
EQ 3 Drums: Toontrack EZdrummer

More realistic effects

Based on the fact that high frequencies are more readily absorbed, reverbs in real rooms contain less high frequencies over time. The frequency content of a reverb also hints at the materials in rooms. Basic equalizers are often built into reverb emulators, or we can apply them externally. We will see later that filters can also reduce the size impression of a reverb.

Stereo enhancements

Frequency differences between the sound arriving at our left and right ears are used by our brain for imaging and localization. Having similar content on the left and right channels, but equalizing each differently widens the perceived image, often creating the impression of a fuller and bigger sound. As already mentioned in Chapter 14, pan-pots only utilize level differences between left and right. Stereo equalization can reinforce the pan-pot function to enhance left/right localization (Figure 15.1).
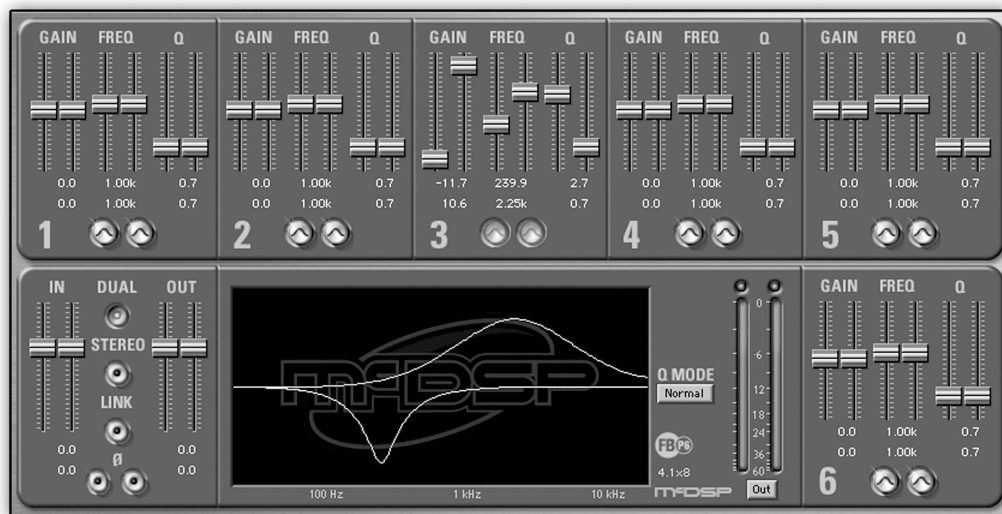
Figure 15.1 Stereo equalization (the McDSP FilterBank P6). A rather unique feature of the FilterBank range is the dual-mono mode (engaged on the left) for stereo inputs. Note that each control has a pair of sliders, each for a different channel. Only band 3 is enabled, and the different settings for each channel are evident in the two curves on the frequency-response display. Such a dual-mono feature enables stereo widening, and more realistic localization.

Fine-level adjustments

Faders boost or attenuate the full frequency spectrum of sounds. Doing so with some instruments (e.g., bass) can alter the overall frequency balance in a mix in a noticeable way. Equalizers let us modify only parts of the spectrum, therefore allowing us finer control over perceived loudness. Many instruments can be made louder by boosting around 3 kHz—our ears' most sensitive area. The advantage in doing this is that other parts of the spectrum remain unchanged, and so masking interaction with other instruments changes less. The disadvantage, however, is that, unlike faders, EQs can negatively change the timbre of the treated instrument.

Better control over dynamic processors

Dynamic range processors are more responsive to low frequencies than high ones. (The reasons for this are explained later in Chapter 17, when we deal with compressors.) Because of this, the lower notes of a bass instrument might be compressed more than the high ones, resulting in inconsistent levels. Or, a kick bleed on a snare track could cause false triggering on a gate. These are just two examples of many problems that an equalizer can rectify when hooked to a dynamic range processor—a combination that also facilitates some advanced mixing techniques. This is expounded upon in the following chapters.

Removing unwanted content

Rumble, hiss, ground loops, mains hum, air-conditioning, the buzz of a light dimmer, and many other extraneous noises can be imprinted on a recording. Spill is also a type of sometimes unwanted content, which gates cannot always remove. For example,

gating is rarely an option when we want to remove the snare bleed from underneath the ride. Equalizers let us filter these types of unwanted sounds.

Compensating for a bad recording

The recording industry would probably have survived had a mysterious virus struck all equalizers and knocked them out for a decade or so—we would simply have had to spend more time miking and use better instruments in better studios. In the meantime, not all recordings are immaculate, so we are happy that our equalizers are healthy.

The frequency spectrum

Humans are capable of hearing frequencies between approximately 20 Hz and 20 kHz (20,000 Hz), though not all people are capable of hearing frequencies up to 20 kHz and our ability to hear high frequencies decreases with age. There are more than a few audio engineers who cannot hear frequencies above 16 kHz, although this does not seem to affect the quality of their mixes or masters.

One very important aspect of our frequency perception is that it is not linear; that is, 100–200 Hz is not perceived equally to 200–300 Hz. Our ears identify with the most fundamental pitch relationship—the octave. We recognize an octave change whenever a frequency is doubled or halved. For instance, 110 Hz is an octave below 220 Hz; 220 Hz is an octave below 440 Hz. While the frequency bandwidth between 110 and 220 Hz is an octave, an identical frequency bandwidth between 10,110 and 10,220 Hz is not even a semitone. Mixing-wise, attenuating 6 dB between 10,110 and 10,220 Hz has little effect compared with attenuating 6 dB between 110 and 220 Hz. As we shall soon see, all equalizers are built to take this phenomenon into account.

Our audible frequency range covers nearly 10 octaves, with the first octave between 20 and 40 Hz and the last octave between 10,000 and 20,000 Hz. Octave-division, how-ever, is commonly done by continuously halving 16 kHz. Many audio engineers develop the skill to identify the 10 different octave ranges, which most of us are able to do with a bit of practice. Keen engineers practice until they can identify ⅓ octave ranges. Although ear-training is beneficial in any area of mixing, it is almost a must when it comes to frequencies. A few gifted people have perfect pitch—the ability to identify (without any reference) a semitone, and even finer frequency divisions.

While engineers talk in frequencies, musicians talk in notes. Apart from the fact that we often need to communicate with musicians, it is beneficial to know the relationship between frequencies and notes. For example, it is useful to know that the lowest E on a bass guitar is 41 Hz, while the lowest E on a standard guitar is an octave above: 82 Hz. We sometimes set high-pass filters (HPFs) just below these frequencies. Also, fre-quencies of notes related to the key of the song might work better than other fre-quencies. We can also deduce the resonant frequency of an instrument by working out the note that excites its body the most. Perhaps the most important frequencies to remember are 262 Hz (middle C), 440 Hz (A above middle C), and 41 Hz (lowest E on a standard bass guitar). Every doubling or halving of these frequencies will deliver the same notes on a different octave. While these frequencies are all rounded, a full table showing the relationship between frequencies and notes can be found in Appendix B.

A famous debate in audio engineering is whether or not people are capable of hearing frequencies above 20 kHz. Science defines hearing as our ability to discern sound through our auditory system (i.e., through our ear canal, eardrum, and an organ called the cochlea, which converts waves into neural discharges that our brain decodes). To date, there has been no scientific experiment that has proved that we hear frequencies above 20 kHz. There have been many that have proved that we cannot. There are studies showing that transmitting high-level frequencies above 20 kHz via bone conduction (metal rods attached to the skull) can improve sound perception among the hearing impaired. But at the same time, it has been proven that the high-frequency content itself is not discerned. Sonar welding systems utilizing frequencies above 20 kHz (such as those used in eye surgery) can pierce a hole in the body. But this is not considered hearing. In the physiological sense of hearing, whatever testing methods researchers used, no subject could hear frequencies above 20 kHz.

It is well understood, though, that frequencies above 20 kHz can contribute to content below 20 kHz and thus have an effect on what we hear. This can be explained by two phenomena called inter-modulation and beating; both can produce frequencies within the audible range as material above 20 kHz passes through a system. Our ears are no such system, but analog and digital systems are. To give an example, inter-modulation happens in any nonlinear device—essentially, any analog component. There is always at least one nonlinear stage in our signal chain—the speakers. Also, if any harmonics are produced within a digital system (due to clipping, for example), harmonics above the Nyquist frequency (half the sampling rate) are mirrored. For this reason, digital clipping at sample rates of 44.1 kHz can sound harsher than at 88.2 kHz.

Regarding the low limit of our audible range, frequencies below 20 Hz are felt via organ resonance in our bodies. But these are not discerned via our auditory system (they would be equally felt by a deaf person).

To conclude: no research has shown that humans are capable of hearing frequencies below 20 Hz or above 20 kHz, but it is well understood that such frequencies can have an effect on what we hear or feel.

Spectral content

A sine wave is a simple waveform that involves a single frequency. Waveforms such as a square wave or a sawtooth generate a set of frequencies—the lowest frequency is called the *fundamental*, and all other frequencies are called *harmonics*. The fundamental and harmonics have a mathematically defined relationship in both level and frequencies. Figure 15.2 shows the spectral content of a 100 Hz sawtooth.

Waveforms such as sines, square waves, and sawtooths constitute the core building blocks of synthesizers. However, the natural sounds we record have far more complex spectral content, which involves more than just a fundamental and its harmonics. The spectral content of all instruments consists of four components (or partials). Combined together, these four partials constitute half of an instrument's timbre (the other half is the dynamic envelope). These are:

- **Fundamental**—the lowest frequency in a sound. The fundamental defines the pitch of the sound.

- **Harmonics**—frequencies that are an integer multiple of the fundamental. For example, for a fundamental of 100 Hz, the harmonics would be 200, 300, 400, 500 Hz, and so forth. Harmonics are an extremely influential part of sounds—we say the fundamental gives sound its pitch, harmonics give the color.
- **Overtones**—frequencies that are not necessarily an integer multiple of the fundamental. A frequency of 150 Hz, for example, would be an overtone of a 100 Hz fundamental. Instruments such as snares tend to have a very vague pitch but produce lots of noise—a typical sound caused by dominant overtones.
- **Formants**—frequencies caused by physical resonances that do not alter in relation to the pitch being produced. Formants are major contributors to a sonic stamp—for example, our ability to recognize each person's voice.

The spectral content of each instrument ranges from the fundamental frequency up to and beyond 20 kHz (even a kick has some energy above 10 kHz). It is worth mentioning that there might be some content below the fundamental due to body resonance or subharmonics caused by the nonlinearity of recording devices. One remarkable faculty of our brain is its ability to **reconstruct missing fundamentals**—whenever the fundamental or the low-order harmonics are removed, our brain uses the remaining harmonics to deduce what is missing. This is the very reason why we still recognize the lower E on a bass guitar as the lowest note, even if our speakers do not produce its fundamental frequency of 41 Hz. This phenomenon has a very useful effect in mixing and will be discussed in greater detail soon.
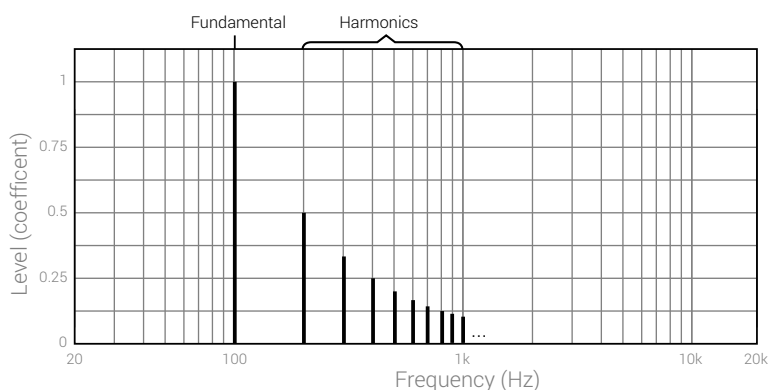
Figure 15.2 The harmonic content of a 100 Hz sawtooth.

Bands and associations

We already know from Chapter 7 that the audible frequency spectrum is divided into four main bands: lows, low-mids, high-mids, and highs. For convenience, the illustration shown in Chapter 7 is reproduced in Figure 15.3. The frequency spectrum can also be subdivided into smaller ranges, each with its own general characteristics. The list below provides very rough guidelines for these different ranges. All the frequencies are approximate and expressed in Hz:

- **Subsonic** (up to 20)—the only instruments that can produce any content in this range are huge pipe organs found in a few churches across the world. This range is
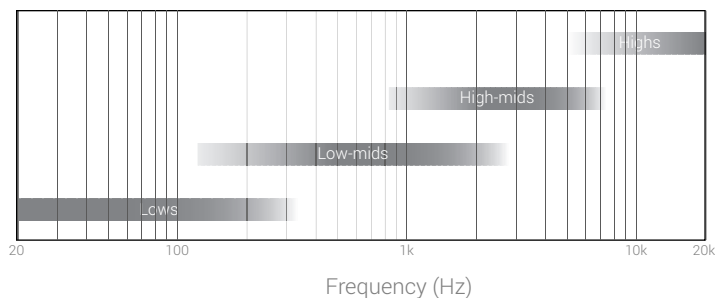
Figure 15.3 The basic four-band division of the audible frequency spectrum.

felt and not heard. Although used in cinemas for explosions and thunder, it is absent from music masters.

- **Low bass** (20–60)—also known as "very lows," this range is **felt** more than heard and is associated with power rather than pitch. The kick and bass usually have their fundamental in this range, which is also used to add sub-bass to a kick. A piano also produces some frequencies in this range.
- **Mid bass** (60–120)—within this range, we start to perceive tonality. Also associated with **power**, mostly that of the bass and kick.
- **Upper bass** (120–250)—most instruments have their **fundamentals** within this range. This is where we can alter the **natural tone** of instruments.
- **Low-mids** (250–2,000)—mostly contain the very important **low-order harmonics** of various instruments, thus their **meat**, **color**, and a big part of their **timbre**.
- **High-mids** (2,000–6,000)—our ears are very sensitive to this range (as per the equal-loudness curves), which contains **complex harmonics**. Linked to **loudness**, **definition**, **presence**, and **speech intelligibility**.
- **Highs** (6,000–20,000)—contain little energy for most instruments, but an important range all the same. Associated with **brilliance**, **sparkle**, and **air**.

The terms in the list above are just a few of many we associate with the various frequency ranges. We also have terms to describe deficiency and excess of various ranges. We use these terms in verbal communication, but we might also use them in our heads—first, we decide that we want to add spark and then we translate it to a specific frequency range. These terms are not standardized, and different people might have different ideas about particular terms. One thing is certain—the frequency ranges we associate these terms with are very rough. The body of a bass guitar is very different from the body of a flute. Figure 15.4 lists these terms.

Sibilance

Sibilance is the whistling sound of s, sh, ch, and sometimes t. "Locks in the castle" has two potentially sibilant consonants. Languages such as German and Spanish have more sibilant constants than English. The same sentence in German—"schlösser im schloss"—has four potentially sibilant consonants. Sibilance is normally found between 2 and 10 kHz and can be emphasized by certain microphones, tube equipment, or just standard mixing tools such as equalizers and compressors. Emphasized sibilance pierces through the mix in an unpleasant and noticeable way. It would also distort on radio transmission and when cut to vinyl.

It is important to note that not all speakers produce the same degree of sibilance. One must be familiar with one's monitors to know when sibilance has to be treated and to what extent. Mastering engineers would probably agree that too many mixes require sibilance treatment. So, if in doubt when mixing, better safe than sorry.

Track 15.2: Sibilance
Since the original vocal recording was not sibilant, an EQ has been employed to draw some sibilance from the vocal. The sibilance is mainly noticeable on "**ci**rcles," "ju**st**," and "pa**ss**." Some untreated vocal recordings can have a much more profound sibilance than the one heard on this track.

Types and controls

Filters, equalizers, and bands

Faders attenuate or boost the whole signal. Put another way, the whole frequency spectrum is made softer or louder. Equalizers attenuate or boost only a specific range of frequencies, and by that they alter the tone of the signal. Early equalizers could only attenuate (filter) frequencies; later designs could boost as well. Regardless, all equalizers are based on a concept known as filtering; thus, the terms *equalizer* and *filter* are used interchangeably. In this text, we will use the term *filter* to mean a circuit that acts with reference to a single frequency, and *equalizer* for a device that might consist of a few filters.

A filter might be in charge of a specific range of frequencies known as *band*. A typical equalizer on a large-format console has four bands: LF, LMF, HMF, and HF (low frequencies, low-mid frequencies, high-mid frequencies, and high frequencies). There would normally also be a high-pass filter (HPF) and a low-pass filter (LPF). Such an equalizer is shown in Figure 15.5. Many software plugins and digital equalizers provide more than four bands, with some plugins offering ten or more. Such equalizers are known as *paragraphic* EQs, and are a hybrid between parametric and graphic equalizers (both explained soon). In many cases, a graph on the plugin window shows the combined frequency response of all bands. A paragraphic EQ plugin is shown in Figure 15.6.

This is probably the right time to look at the difference between the two types of equalizers shown in Figures 15.5 and 15.6. It can be argued that we rarely need more than what is offered by the analog equalizer shown in Figure 15.5—it provides two pass filters to remove content, two shelving filters to change the general tonality of sounds, and two parametric equalizers for more focused treatment. Essentially, most productions before the DAW age were mixed with equalizers that were no more lavish than the one above. Software plugins with a multitude of bands, say ten, can fool people into thinking that standard equalization requires that many bands. It is true that sometimes many bands are useful, but this is more often the case with problematic recordings. For general tone shaping, it is often seasoned engineers that have sufficient auditory skill to take advantage of such a magnitude of bands. Equalization is not the easiest aspect of
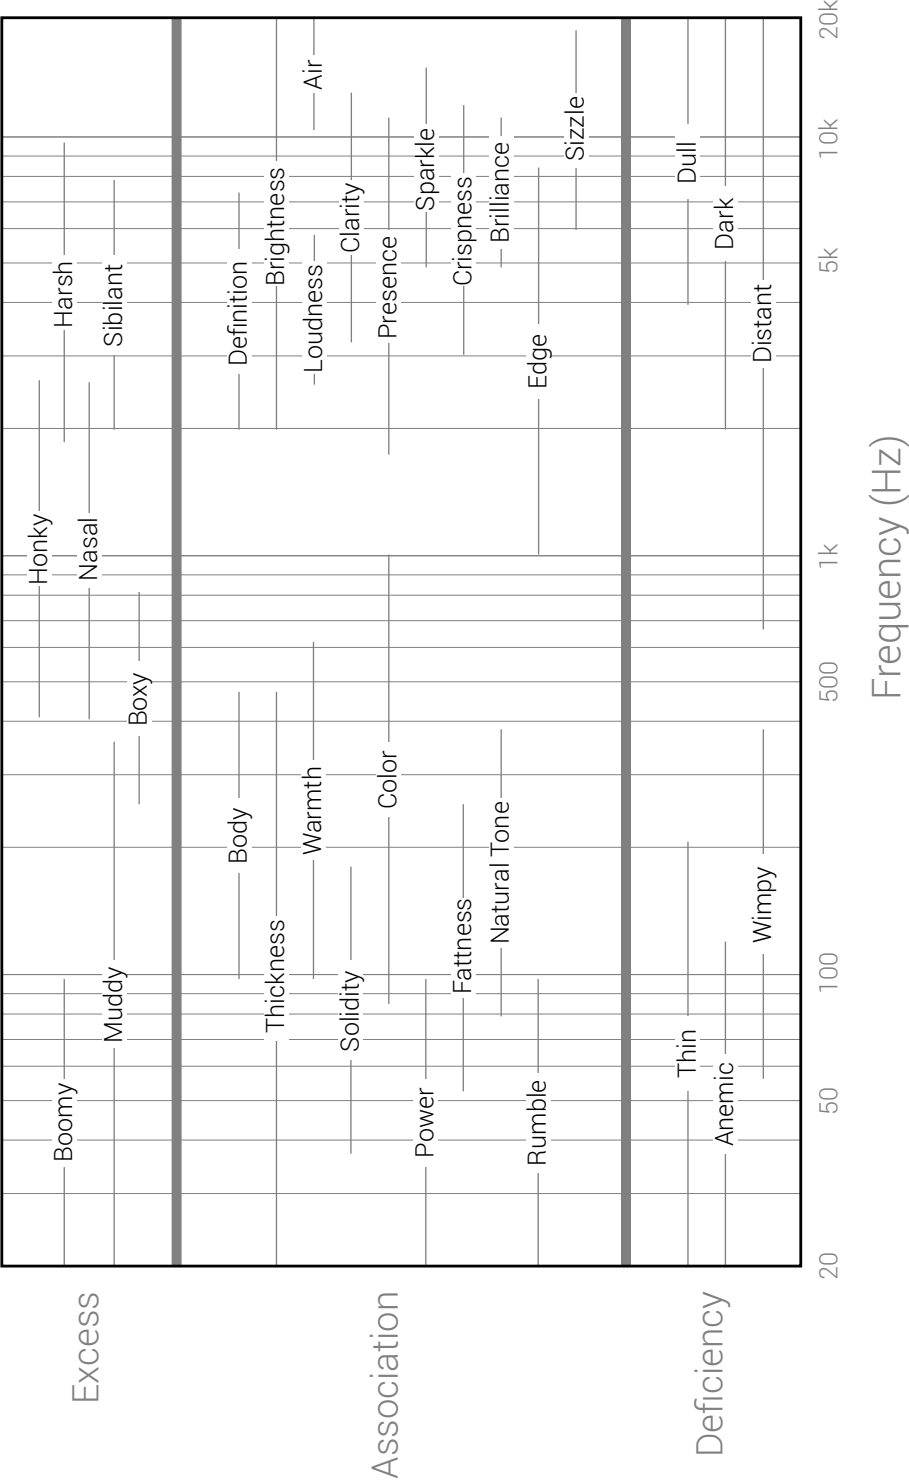
Figure 15.4 Subjective terms we associate various frequency ranges with, and excess or deficiencies in these ranges. The terms are not standardized, and the frequency ranges are rough.
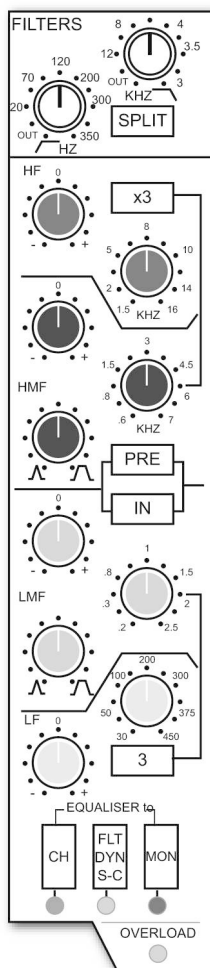
Figure 15.5 A diagram of the SSL 4000 G+ equalizer. The equalizer section consists of four bands: the LF and HF bands are shelving filters; the LMF and HMF are parametric equalizers. At the top are high- and low-pass filters.

mixing, and introducing more bands can make it harder. It is better to hit the right spot with one band than miss it with three.

Another potential problem with paragraphic plugins is the response plot they present. These can easily divert the attention from the actual sonic effect of equalization and lead people to equalize by sight alone. For example, one might over-boost on an equalizer simply because a 6 dB boost on its frequency plot looks graphically small. If you have ever boosted on a paragraphic plugin and then looked at the gain value to see why the effect is so extreme or little, you've probably been equalizing by sight. Professional engineers working on analog desks hardly ever use their eyes while equalizing—you will notice that the gain scales in Figure 15.5 are not even labeled. McDSP lets the user choose whether or not to view the response plot, promoting hearing-based equalization for those who fancy it. It might be wise for other manufacturers to follow suit.
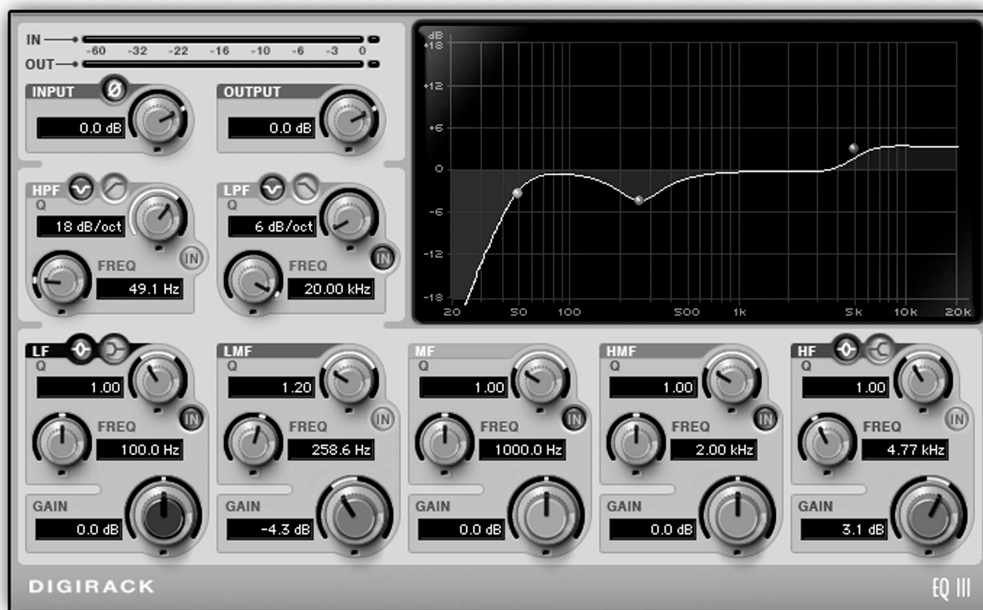
Figure 15.6 A paragraphic equalizer plugin (Digidesign's *DigiRack 7-band EQ 3*). This plugin provides seven bands. Two are high and low filters, and the other five are parametric filters, although the LF and HF bands can be switched into shelving mode.

Frequency-response graphs

We use frequency-response graphs to demonstrate how a device alters the different frequencies of the signal passing through it (Figure 15.7). On these gain vs. frequency graphs, the frequency axis covers the audible frequency range between 20 Hz and 20 kHz. The fact that our perception of pitch is not linear is evident on the frequency scale: 100–200 Hz has the same width as 1–2 kHz. The spacing between the grid lines corresponds to their range (10 Hz steps for tens of hertz, 100 Hz for hundreds of hertz, and so forth).

The filter in Figure 15.7 is known as a *brick-wall filter*. It has a step-like response that divides the frequency spectrum into two distinct bands, where all frequencies to one side are removed and all frequencies to the other are retained. A filter with such a vertical slope is hypothetical—it cannot be built, and even if it could be it would bring about many unwanted side effects. Yet, the term *brick-wall filter* is sometimes used to describe a digital filter with a very steep slope, although never vertical.

There are so many terms and circuit designs involved in filters that it can be very hard to keep track of them all. To generalize, in mixing we use the following types: **pass**, **shelving**, and **parametric**. Each type has a recognizable shape when shown on a frequency-response graph. Regardless of the filter we use, there is always a single reference frequency (e.g., 800 Hz in Figure 15.7), and in most cases we have control over it. This frequency has a different name for each type of filter.
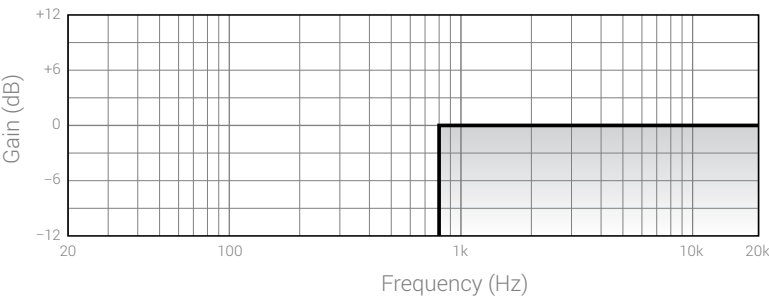
Figure 15.7 A brick-wall filter on a frequency-response graph. The filter in this graph removes all frequencies below 800 Hz, but lets through all the frequencies above. This type of brick-wall filter only exists in theory.

Pass filters

The circuitry of a pass filter can be as simple as a mere combination between a capacitor and a resistor. The reference frequency of pass filters is called the *cut-off frequency*. Pass filters allow frequencies to one side of the cut-off frequency to pass, while continuously attenuating frequencies to the other side. A *high-pass filter* (HPF) allows frequencies higher than the cut-off frequency to pass, but filters frequencies below it. A *low-pass filter* (LPF) does the opposite—it lets what is below the cut-off frequency through, while filtering what is above it. Figure 15.8 illustrates this.

> An HPF can also be referred to as a low-cut filter, and an LPF as a high-cut filter. It is easier to talk about pass filters (rather than cut filters), and more common to use the HPF and LPF abbreviations. The abbreviations LCF and HCF are uncommon in the field of mixing.

It is easy to notice that the cut-off points in Figure 15.8 are not where the curve starts to bend (the transition frequency). Indeed, the cut-off frequency on pass filters is where 3 dB of attenuation occurs. For example, it is 100 Hz for the HPF in Figure 15.8. As a consequence, we can see that a short range of frequencies has been affected despite being higher than the cut-off frequency (or lower, in the case of LPF).

Some filters have a fixed cut-off frequency and only provide an in/out switch. In many cases, the fixed frequency of an HPF would be around 80 Hz, which is right below the
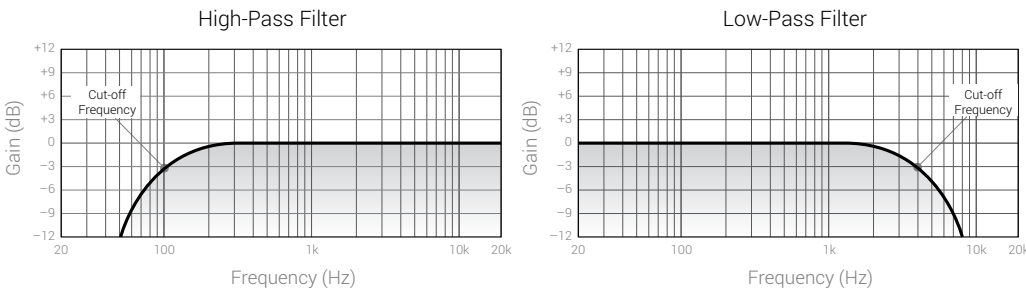
Figure 15.8 A high-pass and a low-pass filter.

lowest note of a regular guitar (E, 82 Hz) and the second harmonic of the lowest note on a bass guitar (E, 41 Hz). However, the majority of the pass filters we use in mixing let us sweep the cut-off frequency.

Another characteristic of a pass filter is *slope*. This determines the steepness of the filter curve. The slope is expressed in dB per octave (dB/oct or dB/8ve), with common values being a multiple of six. A gentle slope of 6 dB/oct means that below the cut-off frequency, each consecutive octave experiences an additional 6 dB of gain loss. With an aggressive slope of 30 dB/oct, it will only take two octaves before frequencies are attenuated by more than 60 dB (a point that can be perceived as effective muting). Generally, the steeper the slope, the more unwanted effects the filter produces. Figure 15.9 shows different slopes. The 6 dB multipliers are set in stone in analog EQs, and rarely can we alter the slope of a filter. An analog pass filter often has a fixed slope of either 6, 12, or 18 dB/oct (in order of popularity). A slope of 36 dB/oct is considered very steep. It is easier in digital designs to offer a variety of slopes, and they are not always bound to 6 dB/oct steps. It should be mentioned that if we do not have slope control, we could achieve our desired slope by combining two filters. For example, we can achieve a 12 dB/oct response by combining two 6 dB/oct filters with the same cut-off frequency. As the signal travels through the first filter, the first octave experiences maximum attenuation of 6 dB; as it travels through the second filter, the same octave experiences an additional 6 dB of attenuation, resulting in a summed response of 12 dB/oct. In the analog domain, this involves connecting two pass filters in a series. With a paragraphic plugin, this involves having two bands set to the same pass response. In all cases, the cut-off frequencies of the two filters should be identical.

Although this might seem obvious, a few people wrongly assume that once a signal is passed through a filter, a second pass through the same filter will have no effect. Filters always attenuate and never remove completely. With a pass filter, for example, during the first signal pass the cut-off frequency will be attenuated by 3 dB, then during the second pass it will be attenuated by an additional 3 dB.

The following tracks demonstrate the effect of an LPF and an HPF on drums. The cut-off frequency used in each sample is denoted in the track name. All slopes are 24 dB/oct.

Track 15.3: Drums Source
The source, unprocessed track is used in the following samples:

Track 15.4: HPF 50 Hz (Drums)

Track 15.5: HPF 100 Hz (Drums)

Track 15.6: HPF 250 Hz (Drums)

Track 15.7: HPF 2 kHz (Drums)

Track 15.8: HPF 6 kHz (Drums)

Track 15.9: LPF 12 kHz (Drums)

Track 15.10: LPF 6 kHz (Drums)

Track 15.11: LPF 2 kHz (Drums)

Track 15.12: LPF 250 Hz (Drums)

Track 15.13: LPF 100 Hz (Drums)

And the same set of samples with vocal:

Track 15.14: Vocal Source

Track 15.15: HPF 50 Hz (Vocal)

Track 15.16: HPF 100 Hz (Vocal)

Track 15.17: HPF 250 Hz (Vocal)

Track 15.18: HPF 2 kHz (Vocal)

Track 15.19: HPF 6 kHz (Vocal)

Track 15.20: LPF 12 kHz (Vocal)

Track 15.21: LPF 6 kHz (Vocal)

Track 15.22: LPF 2 kHz (Vocal)

Track 15.23: LPF 250 Hz (Vocal)

Track 15.24: LPF 100 Hz (Vocal)

Plugin: McDSP FilterBank F2
Drums: Toontrack EZdrummer

The following tracks demonstrate the effect of different filter slopes on drums and vocals. The cut-off frequency and slope used in each sample are denoted in each track name (slopes are in dB/oct).

Track 15.25: HPF 250 Hz Slope 6 (Drums)

Track 15.26: HPF 250 Hz Slope 12 (Drums)

Track 15.27: HPF 250 Hz Slope 18 (Drums)

Track 15.28: HPF 250 Hz Slope 24 (Drums)

Track 15.29: LPF 6 kHz Slope 6 (Drums)

Track 15.30: LPF 6 kHz Slope 24 (Drums)

Track 15.31: HPF 250 Hz Slope 6 (Vocal)

Track 15.32: HPF 250 Hz Slope 12 (Vocal)

Track 15.33: HPF 250 Hz Slope 18 (Vocal)

Track 15.34: HPF 250 Hz Slope 24 (Vocal)

Track 15.35: LPF 6 kHz Slope 6 (Vocal)

Track 15.36: LPF 6 kHz Slope 24 (Vocal)

Plugin: McDSP FilterBank F2
Drums: Toontrack EZdrummer

Figure 15.9 Different slopes on an HPF. These four instances of the Digidesign *DigiRack EQ 3* illustrate a different slope each. From top to bottom: 6, 12, 18, and 24 dB/oct. The cut-off frequency is set to 1 kHz and is clearly indicated with a white circle.

Many of us are familiar with the pass filters on synthesizers that have both cut-off and resonance control. Resonance provides a boost around the cut-off frequency, and gives some added edge to the transition range. Resonance is highly noticeable when the cutoff frequency is swept, and most DJ mixers incorporate resonant filters. Resonance is not an extremely common feature in mixing equalizers, although some, such as the PSP *MasterQ* in Figure 15.10, do provide it. However, it might still be "secretly" incorporated into some designs, often those that offer very few controls and do not reveal their frequency-response graph (typically analog or the digital plugins that emulate them). One such example is the Universal Audio *NEVE 1073 EQ*, shown in Figure 15.11. Generally speaking, equalizers of this kind have a frequency response that is far from textbook perfect (such as Figure 15.8), a fact that often contributes to their appealing sound. Even if an equalizer
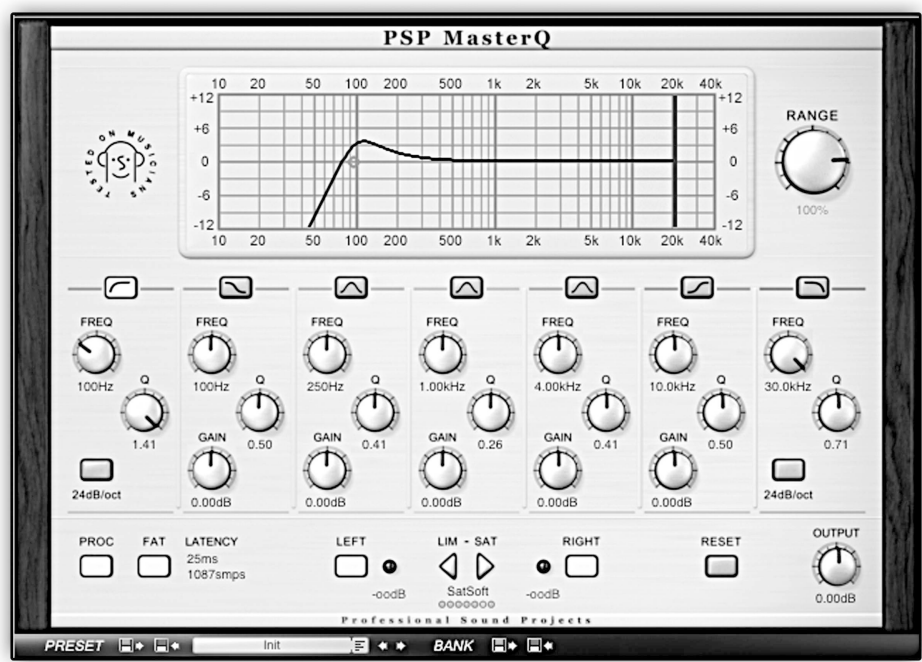
Figure 15.10 Pass filter resonance (the PSP MasterQ plugin). The Q control on the HPF section acts as a resonance control. The resonance is seen as a bump around the cut-off frequency.

Figure 15.11 The Universal Audio NEVE 1073 EQ plugin. This plugin, which emulates the sound of the legendary analog NEVE 1073, has a frequency response that deviates from the perfect theoretical shape of filters, a fact that contributes to its appealing sound. The high-pass filter (rightmost control) involves changing the resonance for each of the four selectable cut-off frequencies.
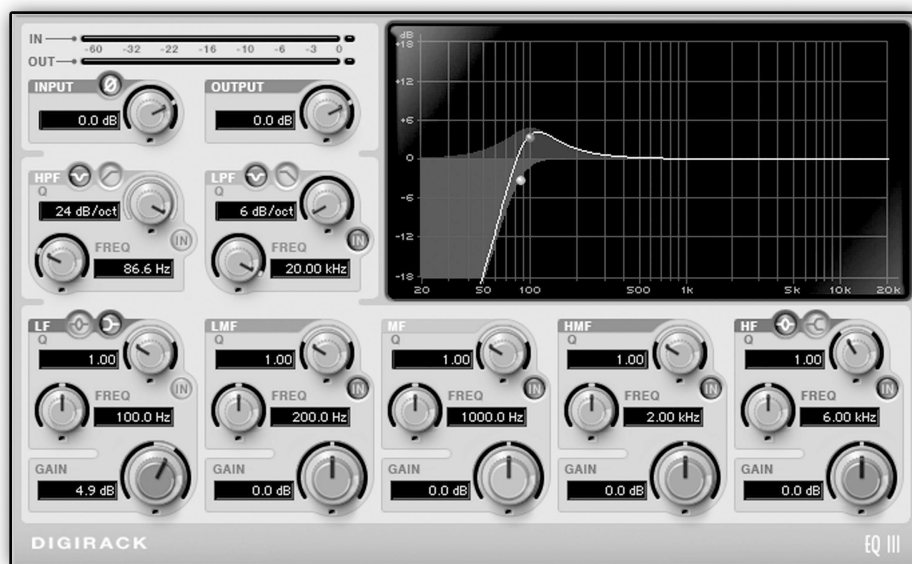
Figure 15.12 Combining two bands to create a resonant filter (the Digidesign *DigiRack 7-band EQ 3*). A parametric filter (LF band) is used to create a bump around the cut-off frequency of an HPF, resulting in a response typical of a resonant filter.

does not offer a resonance control, we can achieve this characteristic by adding a parametric filter around the cut-off frequency. Figure 15.12 illustrates how this is done. The only limitation with such a setup is that both reference frequencies have to be adjusted if we want to sweep the resultant response.

Track 15.37: HPF Sweep No Resonance
AN HPF set to 250 Hz with no resonance starts sweeping up in frequency after the second bar.

Track 15.38: HPF Sweep Resonance
Similar arrangement as in the previous track, only with resonance.

Plugin: PSP MasterQ
Drums: Toontrack EZdrummer

Shelving filters

Most people have used shelving filters—these are the bass and treble controls found in our domestic hi-fi systems, also known as tone controls. Shelving filters, as they are now known, were conceived by Peter Baxandall in the late 1940s. They are so-called as their response curve can, in inspiring moments at least, remind us of shelves.

Unlike pass filters, which only cut frequencies, shelving filters can also boost. The reference frequency of shelving filters divides the spectrum into two bands. On one
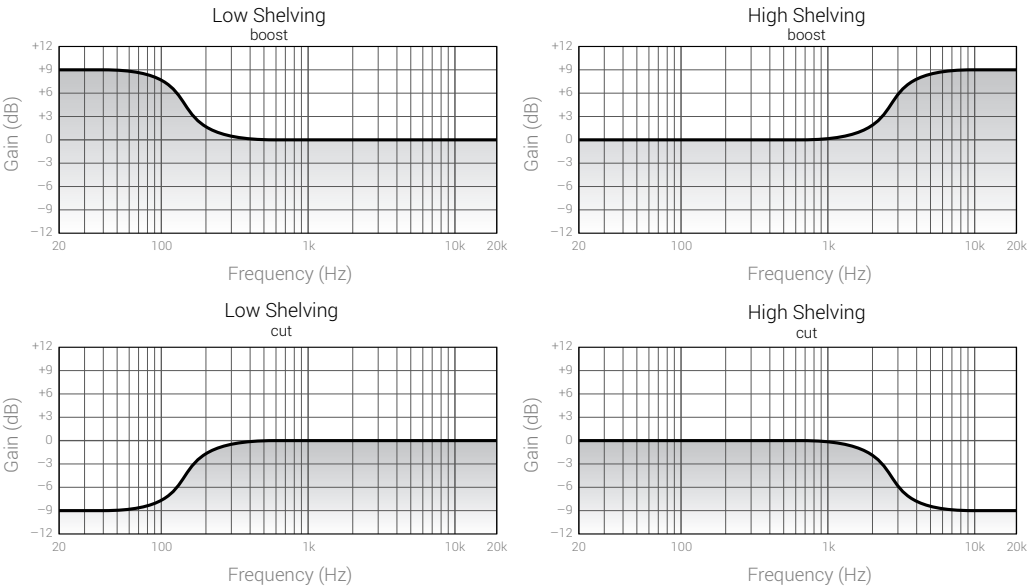
Figure 15.13 The four versions of shelving filters. For boost, +9 dB of gain is applied; for attenuation, –9 dB.

side, frequencies are undisturbed; on the other, frequencies are either attenuated or boosted by a constant amount. A gain control determines that amount. As per the above discussion about vertical response slopes being impossible, there is always a transition band between the unaffected frequencies and those affected by the set gain. Figure 15.13 shows the four possible versions of shelving filters.

When it comes to defining the reference frequency of shelving filters, we encounter ambiguity. Designers might choose one of the three main possibilities (illustrated in Figure 15.14). Some define it in the traditional engineering sense as the point at which 3 dB of gain is reached—the familiar *cut-off frequency*. However, this tells us little
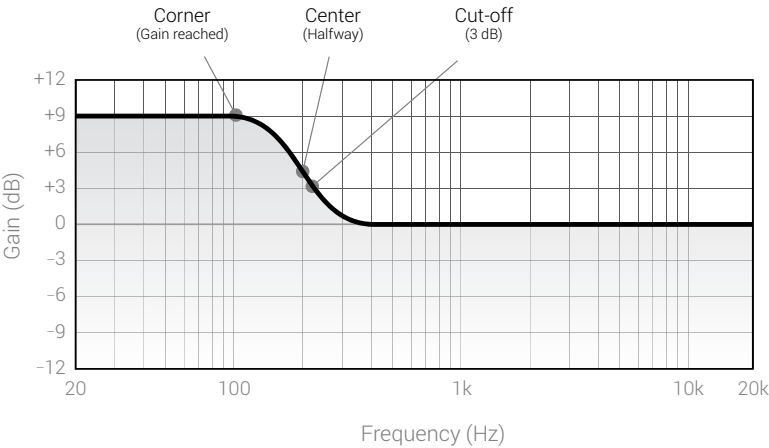
Figure 15.14 Three possible options for the shelving frequency. The corner frequency is where the set gain is reached. The center frequency is halfway through the transition band. The cut-off frequency is the traditional 3 dB point.

about the real effect of the filter, which roughly happens when the set amount of gain is reached. Since this is often what we are after, some designers use this point as a reference, and it is called the *corner frequency*. To add to the confusion, it is also possible for the reference frequency to be halfway on the transition range—a frequency we can regard as the *center frequency*. Of all three options, it can be argued that the corner frequency is the most intuitive one to work with.

All shelving filters offer control over the gain amount, often ranging between −12 and +12 dB. Most filters in mixing also offer control over the shelving frequency. Some filters also offer slope control, which determines the steepness of the slope in the transition band. Just like with pass filters, the actual response of shelving filters might deviate from the curves shown in Figure 15.13. In fact, it is likely to deviate. One very common response involves a contrast resonance around the transition frequency; that is, a section of the curve that bends in the opposite direction to the normal response. Such a response can be seen in Figure 15.15.
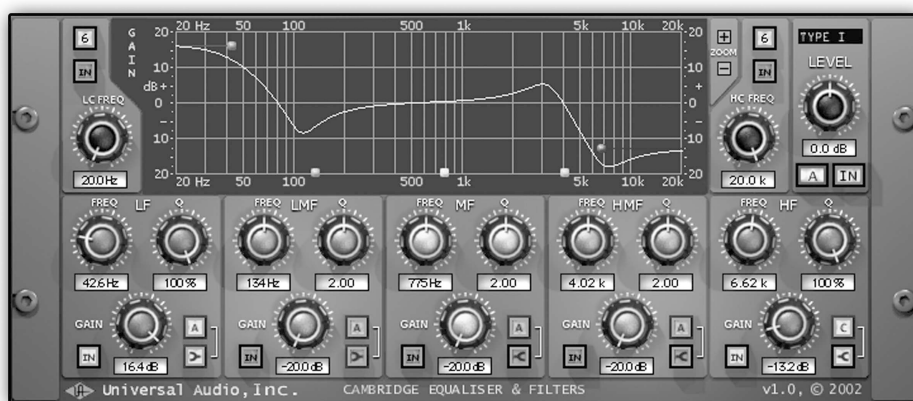
Figure 15.15 A typical contrast resonance on a shelving filter (the Universal Audio *Cambridge EQ*). Both the low- and high-shelving filters are engaged in this screenshot, and it might be easier to discern them if we imagine a cross-line at 700 Hz. The low-shelving filter, set to type A, has a single contrast bend around the transition frequency. The high-shelving filter, set to type C, has two of these bends—one around the transition frequency, the other around the corner frequency.

The following tracks demonstrate the effect of different boost and attenuation amounts on low- and high-shelving filters (LSF and HSF). The cut-off frequency and amount of gain used in each sample are denoted in each track name.

Track 15.39: LSF 250 Hz 3 dB Down (Drums)

Track 15.40: LSF 250 Hz 6 dB Down (Drums)

Track 15.41: LSF 250 Hz 12 dB Down (Drums)

Track 15.42: LSF 250 Hz 20 dB Down (Drums)

Track 15.43: LSF 250 Hz 3 dB Up (Drums)

Track 15.44: LSF 250 Hz 9 dB Up (Drums)

Track 15.45: HSF 6 kHz 3 dB Down (Drums)

Track 15.46: HSF 6 kHz 9 dB Down (Drums)

Track 15.47: HSF 6 kHz 20 dB Down (Drums)

Track 15.48: HSF 6 kHz 3 dB Up (Drums)

Track 15.49: HSF 6 kHz 9 dB Up (Drums)

Track 15.50: LSF 250 Hz 3 dB Down (Vocal)

Track 15.51: LSF 250 Hz 6 dB Down (Vocal)

Track 15.52: LSF 250 Hz 12 dB Down (Vocal)

Track 15.53: LSF 250 Hz 20 dB Down (Vocal)

Track 15.54: LSF 250 Hz 3 dB Up (Vocal)

Track 15.55: LSF 250 Hz 9 dB Up (Vocal)

Track 15.56: HSF 6 kHz 3 dB Down (Vocal)

Track 15.57: HSF 6 kHz 9 dB Down (Vocal)

Track 15.58: HSF 6 kHz 20 dB Down (Vocal)

Track 15.59: HSF 6 kHz 3 dB Up (Vocal)

Track 15.60: HSF 6 kHz 9 dB Up (Vocal)

Plugin: Sonnox Oxford EQ+Filters
Drums: Toontrack EZdrummer

Parametric filters

In 1972 at the AES, George Massenburg unveiled the parametric equalizer—a revolutionary circuit that he designed with help from fellow engineers. Although the concept of bandpass and band-reject filters (primitive types of parametric filters) was already well established, parametric equalizers became, and still are, de facto in mixing.

Parametric filters can cut or boost. Their response curve is reminiscent of the shape of a bell, as can be seen in Figure 15.16. The reference frequency is called the *center frequency* and we can sweep it higher or lower. The gain determines the maximum amount of boost or cut reached at the center frequency. The two cut-off points are 3 dB away from the center frequency (3 dB below for boost, 3 dB above for cut). The bandwidth is measured between these two cut-off points, and we express it in octaves.

Had we expressed the bandwidth in hertz (for example, 600 Hz for the graphs in Figure 15.16), the effect of the filter would alter as the center frequency is swept, where the higher the frequency is, the less the effect. Consequently, the bell shape would narrow as the center frequency is swept higher. The reason for this has to do with our nonlinear pitch perception. To demonstrate this again, 600 Hz between 200 and 800 Hz equals two octaves (24 semitones); the same 600 Hz between 10,000 and 10,600 Hz is only a semitone. There is no comparison between affecting two octaves and a semitone.
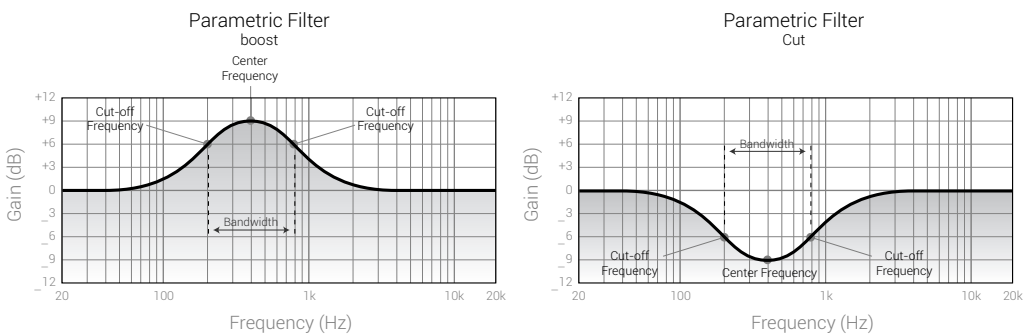
Figure 15.16 A parametric filter. Both response graphs involve 9 dB of gain (boost or cut) and a center frequency at 400 Hz; 3 dB below the center frequency are the two cut-off frequencies at 200 and 800 Hz. The two-octave bandwidth (200–800 Hz) is measured between the two cut-off points.

Although the bandwidth on some equalizers is expressed in octaves, it is far more common to use a parameter called *Q* (*quality factor*). *Q* can be calculated by the mathematical expression *Fc* / (*Fh* – *Fl*), where *Fc* is the center frequency and *Fh* and *Fl* represent the high and low cut-off frequencies, respectively. The higher the *Q*, the narrower the shape of the bell. Roughly speaking, *Q* values range from 0.1 (very wide) to 16 (very narrow). Three different *Q* settings can be seen in Figure 15.17.

The shape of the bell gives the filter much of its characteristics, and it is not surprising that many variants exist. One important aspect is whether or not there is a dependency between the gain and the *Q* (Figure 15.18). With some designs, the bell narrows with gain (a behavior described as *proportional Q*). As a consequence, changing the gain might also require adjustment to the *Q*. Equalizers of this type tend to sound more
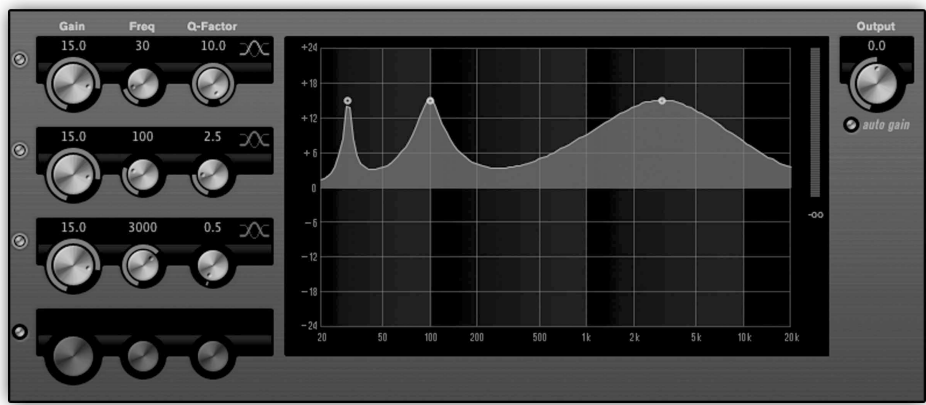
Figure 15.17 Different *Q* settings (the Cubase *StudioEQ*). Three bands are engaged in this screenshot, all with a gain boost of 15 dB. The lowest band (leftmost) shows a response with a narrow *Q* (10). The middle band response is achieved with a moderate *Q* (2.5). The widest *Q* (0.5) is applied on the highest band. The different bandwidths can be visualized by looking at the +12 dB grid line between the cut-off points.
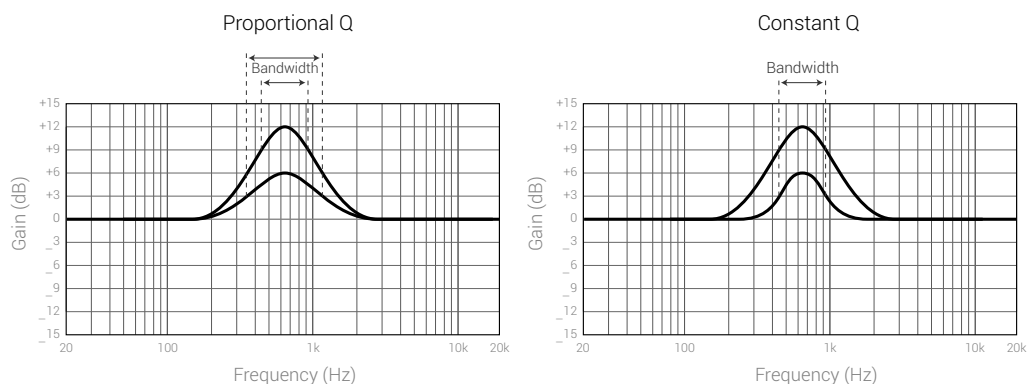
Figure 15.18 Proportional vs. constant Q. With proportional Q, the band-width varies in relation to the gain; with constant Q, it remains the same.

forceful as they become sharper with higher gain settings. A design known as *constant Q* provides an alternative where the bandwidth is (nearly) constant for all gain settings. This produces a softer effect that often brings about more musical results.

In this book, the term wide Q *denotes a wide bell response that is achieved using low* Q *settings (such as 0.1). The term* narrow Q *denotes a narrow response that is the outcome of high* Q *settings (such as 16).*

Track 15.61: Pink Noise Automated *Q*
This sample, which unmistakably resembles the sound of a sea wave, is the outcome of an equalized pink noise. The initial settings include a 9 dB boost at 1 kHz with the narrowest Q of 10. Due to the narrow Q and the large boost, it is possible to hear a 1 kHz whistle at the beginning of this track. In the first 8 seconds, the Q widens to 0.1, a period at which the whistle diminishes and both low and high frequencies progressively become louder. For the next 8 seconds, the Q narrows back to 10, a period when the augmentation of the 1 kHz whistle might become clearer.

Plugin: Digidesign DigiRack EQ 3

Most of the sounds we are mixing have a rich and complex frequency content that mostly ranges from the fundamental frequency to 20 kHz and beyond. The timbre components of various instruments do not exist on a single frequency only, but stretch across a range of frequencies. One of the accepted ideas in equalization is that sounds that focus on a very narrow frequency range are often gremlins that we would like to remove. To accommodate this, designs provide an attribute called *boost/cut asym-metry*, which has grown in popularity in recent years. In essence, an asymmetrical filter of this type will use a wider bell response for boosts, but a narrower one for cuts. We, the users, see no change in the *Q* value. This type of asymmetrical response is illus-trated in Figure 15.19.

Figure 15.19 Cut/boost asymmetry (the Sonnox *Oxford Equalizer and Filters* plugin). Both the LMF and HMF are set with extreme gain of –20 and +20 dB, respectively. The *Q* on both bands is identical (2.83). The asymmetry is evident as the cut response is far narrower than the boost response. This characteristic is attributed to the EQ-style selection seen as Type-2 in the center below the plot. Other EQ styles on this equalizer are symmetrical.

Sounds that focus on a very narrow frequency range are likely to be flaws that we will want to eliminate.

In the following tracks, a parametric filter is applied on drums. The center frequency is set to 1 kHz, the Q to 2.8, and the gain is automated from 0 up to 16 dB and back down to 0. Notice how in the first track, which involves proportional Q, the operation of the filter seems more obstructive and selective, whereas in the second track, which involves constant Q, it seems more subtle and natural.

Track 15.62: Proportional Q

Track 15.63: Constant Q

Plugin: Sonnox Oxford EQ+Filters
Drums: Toontrack EZdrummer

A variation of parametric filter is known as a *notch filter*. We use this term to describe a very narrow cut response, such as the cut in Figure 15.19. This type of response is often used to remove unwanted frequencies, such as mains hum or strong resonance.
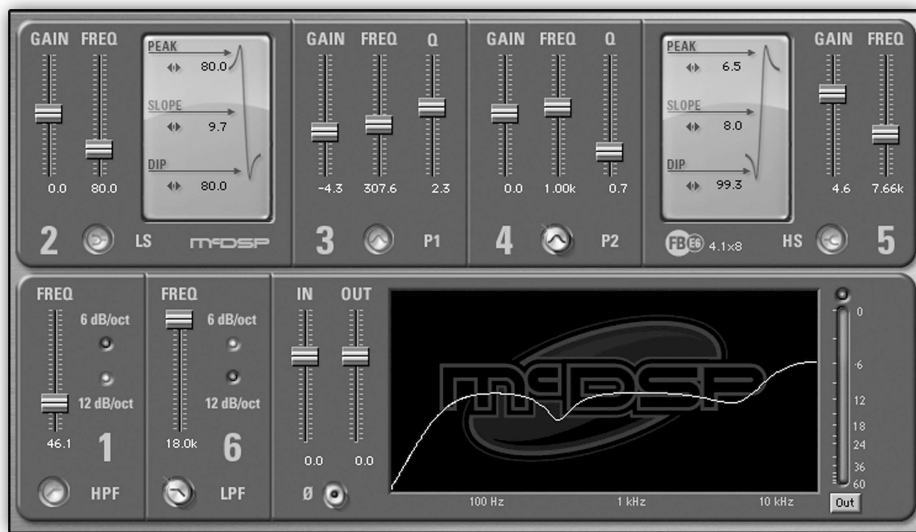
Figure 15.20 The McDSP *FilterBank E6*.

A summary of pass, shelving, and parametric filters

It is worthwhile at this stage to summarize the various filters we have looked at so far. The McDSP plugin in Figure 15.20 will help us recap.

A **pass filter** continuously removes frequencies to one side of the cut-off frequency. Normally we have control over the **cut-off frequency**, and occasionally we can also control the slope. The HPF and LPF are bands 1 and 6, respectively, in Figure 15.20.

A **shelving filter** boosts or attenuates frequencies to one side of the corner frequency. Normally we can control the **corner frequency** (or any other reference frequency if different) and **gain**. The shelving bands in Figure 15.20 are 2 and 5. The *FilterBank E6* also provides control over the peak and dip resonance and the slope.

A **parametric filter** normally provides **gain**, **frequency**, and *Q*. The set gain is reached at the center frequency, and the *Q* (or octave bandwidth) determines the width of the bell between the two cut-off points. Bands 3 and 4 in Figure 15.20 are parametric filters. A parametric filter that offers variable gain, frequency, and *Q* is known as a *(fully) parametric equalizer*. A parametric filter that only offers variable gain and frequency is known as a *semi-parametric* or *sweep EQ*.

The response characteristics of various filter designs give each equalizer its unique sound. Deviations from the textbook-perfect shapes, among other factors, gave many famous analog units their distinctive and beloved sounds. To our delight, it is not uncommon nowadays to come across professional plugins that provide different types of characteristics to choose from and more control over the equalizer response.

Graphic equalizers

A graphic equalizer (Figure 15.21) consists of many adjacent mini-faders. Each fader controls the gain of a bell-response filter with fixed frequency, acting on a very narrow band. The *Q* of each band is fixed on most graphic equalizers, yet some provide variable *Q*. The frequencies are commonly spaced either an octave apart (so ten faders are used to cover the audible frequency range) or a third of an octave apart (so 27–31 faders are used). Graphic equalizers are so-called because the group of faders gives a rough indication of the applied frequency response.

Graphic equalizers are very common in live sound, where they are used to tune the venue and prevent feedback. However, they are uncommon in mixing due to their inherent limitations compared with parametric equalizers. The multitude of filters involved (up to 31) means that many hardware units compromise on quality for the sake of cost. Arguably, software plugins can easily offer a graphic EQ of better quality than most analog hardware units on the market. But there are not many situations where such a plugin would be favored over a parametric equalizer.

Graphic equalizers are the standard tools in **frequency training**, where the fixed and limited amount of frequencies is actually an advantage. As pink noise is played through the equalizer, one person boosts a random band and another person tries to guess what band has been boosted. The easiest challenge involves pink noise, focusing on a limited number of bands (say eight) and generous boost such as 12 dB. Things get harder with lower gain boosts, cutting instead of boosting, adding more bands, and playing real recordings rather than noise.

We can train our ears even when alone. For example, playing a kick through a graphic equalizer, going octave by octave and attentively listening to how a boost or cut affects the kick's timbre, is a beneficial exercise. It almost goes without saying that this can also be done with a parametric equalizer, but graphical equalizers are tailored to the task. Both the *Marvel* and *Overtone GEQs* by Voxengo are free cross-platform plugins suited for this sort of training.
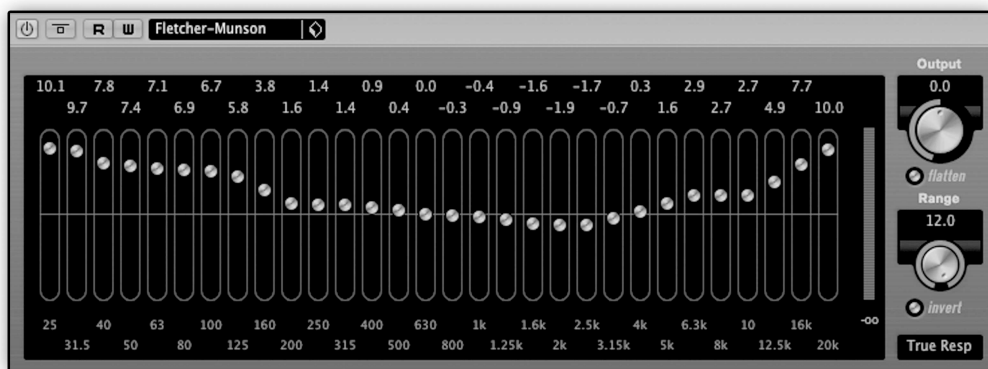
Figure 15.21 A 30-band graphic equalizer plugin (the Cubase *GEQ-30*). The fader settings shown here are the outcome of the Fletcher-Munson preset, which is based on the equal-loudness contours.

Highly trained engineers can identify gain changes as small as 3 dB in ⅓ octave spacing. Trained ears make equalization easier as we can immediately recognize which frequencies need treating. Any masking issues are readily addressed, and we have a much better chance of crafting a rich and balanced frequency spectrum.

Graphic equalizers are great for frequency training.

Dynamic equalizers

Dynamic equalizers are not currently that popular and are often associated with mastering applications. Yet, the plugin revolution means that we should expect to see more of them in the future, and they can be just as beneficial in mixing as they are in mastering.

In contrast to the standard (static) equalizer, where the amount of cut or boost on each band is constant, the same amount on a dynamic equalizer is determined by the gain intensity on each band. Put another way, the louder or softer a specific band becomes, the less or more cut or boost is applied. Dynamic equalizers are something of a marriage between a multiband equalizer and a multiband compressor. For each band, we often get the familiar compressor controls of threshold, ratio, attack, and release. Unlike multiband compressors, these do not control the gain applied on the frequency band, but the amount of boost or cut on the equalizer of that band. Figure 15.22 provides a diagram of a one-band dynamic equalizer, while Figure 15.23 shows a multiband dynamic EQ.

Dynamic equalizers are very useful when we want to correct varying frequency imbalances, such as those caused by the proximity effect. We can set the EQ to cut more lows when these increase in level. We can also reduce finger squeaks on a guitar only when these happen. There are many more fascinating applications, but these are
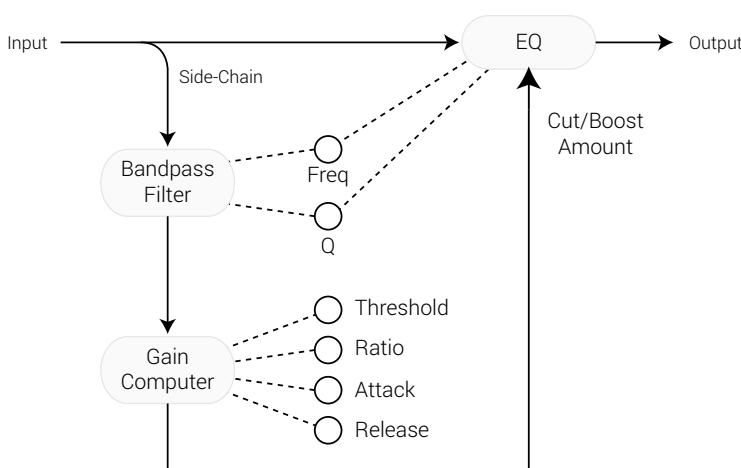
Figure 15.22 Basic diagram of single-band dynamic EQ. The frequency and *Q* settings determine both the equalizer settings and the pass-band frequencies that the gain computer is fed with. Basic compressor controls linked to the gain computer dictate the amount of cut or boost applied on the EQ.
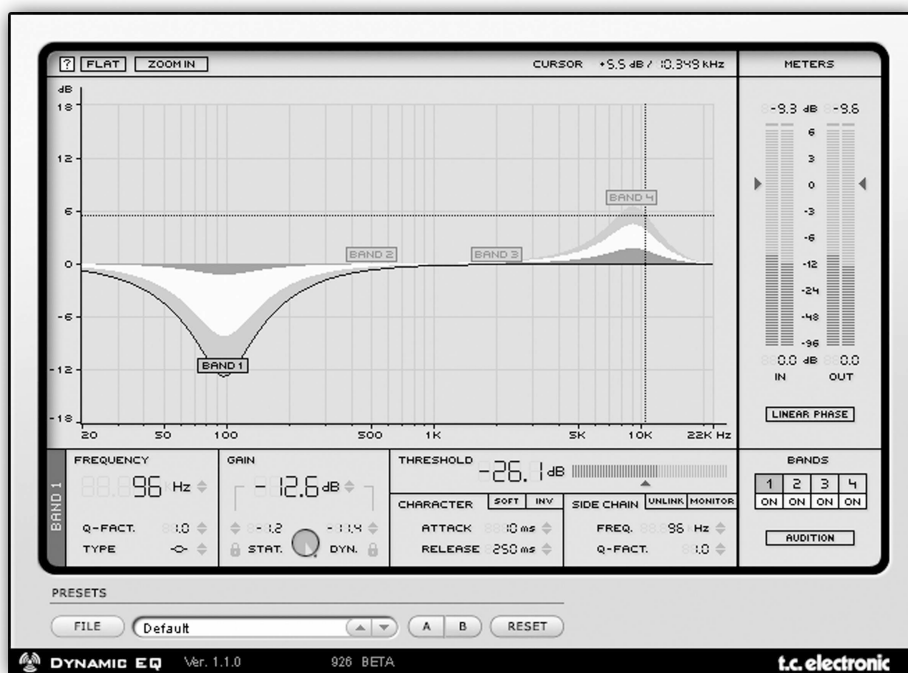
Figure 15.23 Dynamic EQ plugin (the t.c. electronic *Dynamic EQ*). Both bands 1 and 4 are active in this screenshot. Looking at each band, the darkest area (most inward) shows static equalization, the light area shows dynamic equalization, and the gray area (most outward) shows the band-pass filter response curve.

explored in Chapter 17, which is about compressors, as they are more common at present than dynamic equalizers.

In practice

Equalization and solos

Mostly, equalization is done in mix-context. Solving masking or tuning an instrument into the frequency spectrum is done in relation to other instruments. Equalizers and solo buttons are not exactly best friends. Instruments that sound magnificent when soloed can sound awful in the mix. The opposite is also likely—an instrument that sounds dreadful when soloed can sound great in the mix. If we apply high-pass filtering to soloed vocals, we are likely to set a lower cut-off point than we would have had we listened to the whole mix—in isolation, we only hear what the filter removes from the sound, but we cannot hear how this removal can give the vocals more presence and clarity in the mix. Yet, there are several situations where we equalize while soloing. Mostly, this involves initial equalization and instances where the mix clouds the details of our equalization. But it is worth remembering that:

Equalizing a soloed instrument can be counter-effective.

Upward vs. sideways

An old engineering saying is that if a project was recorded properly, with the mix in mind, there would be little or no equalization to apply. To some recording engineers, this is a major challenge. We say once sounds are recorded that they exist in their purest form and any equalization is interference with this purity. In subtle amounts, equalization is like makeup. But in radical amounts, it is like pretty drastic plastic surgery—it can bring dreadful results. I like to compare equalizers to funfair mirrors—gently curved mirrors can make you slimmer, broader, taller, or shorter, which sometimes make you look better. But the heavily curved mirrors just make you look funny and disproportionate. Having said that, as part of the sonic illusion we provide in some mixes, equalizers are used generously for extreme manipulations of sounds. The natural versus artificial concept is very relevant when it comes to equalization.

There is no doubt that some equalizers lend themselves better to our artistic intentions than others, but they all share common issues. The more drastically we employ them, the more we stand to lose in return. In that sense, a perfect EQ is one that has a flat frequency response, and thus no effect at all. In this ideal state, we are more likely to impair the sound with:

- more gain;
- narrower *Q* settings;
- steeper slopes; and
- more angular transition bends.

Still, equalizers are not that injurious—they are widely used in mixes and yield magnificent results. The point is that sometimes we can achieve better results with a slight change of tactic—one that simply involves less drastic settings.

Say, for example, Sid wants to add attack to a kick. Sid reads in some book that the kick's attack is at 6 kHz. So he dials a parametric EQ to that frequency, and the more attack he wants, the more he boosts the gain. Sid equalizes *upward*. He ends up with a narrow *Q* and +12 dB of gain. Nancy knows that very few specific sounds focus on a narrow frequency range, and these are often gremlins that we want to remove, not emphasize. So she dials a frequency of 6 kHz as well, but with a wider *Q* and lower gain. Nancy works *sideways*. Nancy boosts less, but boosts a wider range of the kick's attack, so the overall impact is roughly the same, only with less artifacts. This example is illustrated in Figure 15.24a.

We can use the same example with other types of filters. We can set a high-shelving filter at one frequency with more gain or set it to a lower frequency with lower gain (Figure 15.24b). Similarly, we can set an HPF with a steep slope at one frequency or we can set a gentler slope at a higher frequency (Figure 15.24c). The performance of a pass filter is often evaluated by how well it handles the transition area, especially with steeper slopes. Cheap designs can produce very disturbing side effects around the cut-off frequency. The three examples in Figure 15.24 can be summarized as:

> *Q (for parametric filter) and frequency (for pass or shelving) can be traded for gain or slope, resulting in "cleaner" equalization.*
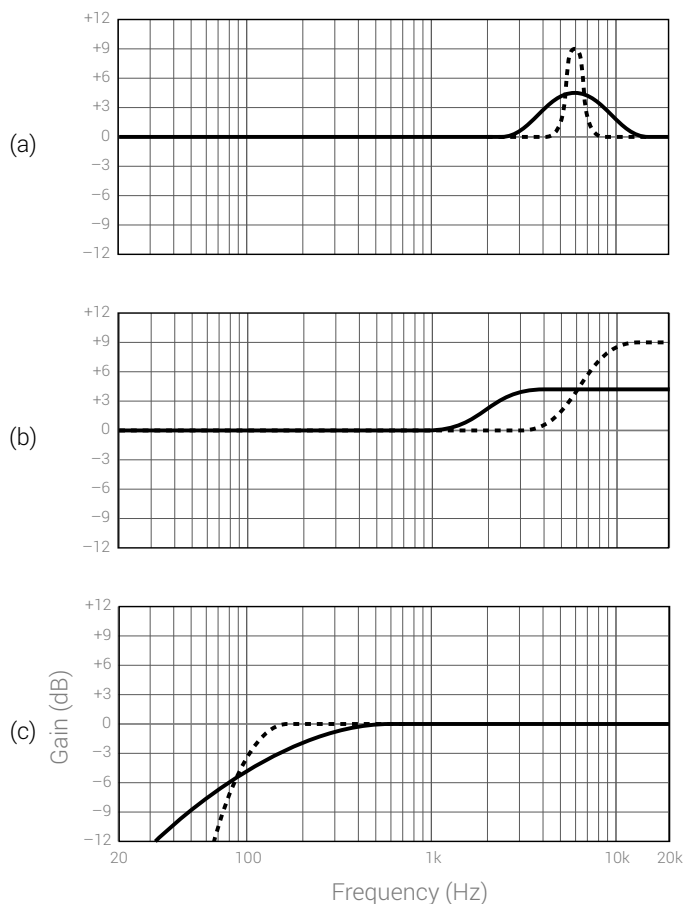
Figure 15.24 Equalization alternatives. In all these graphs, the dashed curves involve more drastic settings than the solid curves. The dashed and solid curves could bring about similar results, although the dashed curves are likely to yield more artifacts. This is demonstrated on a (a) parametric filter, (b) shelving filter, and (c) pass filter.

We must not forget that, as with many other mixing tools, sometimes we are more interested in hearing the edge—subtlety and transparency is not always what we are after. For example, in genres such as death metal, equalizers are often used in what is considered a radical way, with very generous boosts. The equalizer's artifacts are used to produce harshness, which works well in the context of that specific music. Some equalizers have a very characteristic sound that only sharpens with more extreme settings.

Equalizers and phase

The operation of an equalizer involves some delay mechanism. The delays are extremely short and well below 1 ms. *Group delay* is a term often used, suggesting that only specific frequencies are delayed—while not strictly precise, it is faithful to the fact that some frequencies are affected more than others. Regardless, the delay mechanism in equalizers results in unwanted phase interaction. Just like two
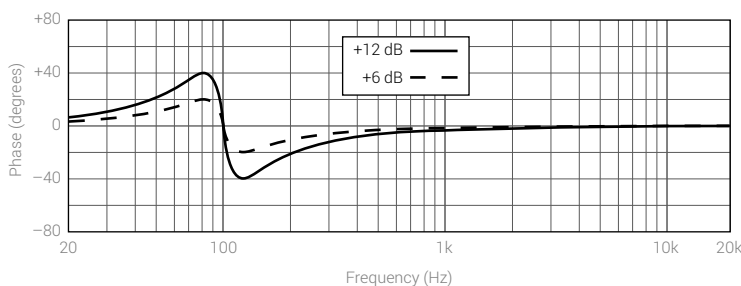
Figure 15.25 The phase response of a boost on a parametric equalizer. The equalizer center frequency is 100 Hz. The solid line shows a boost of 12 dB, and the dashed line shows a gain boost of 6 dB.

identical signals that are out of phase result in comb filtering, we can simplify and say that an equalizer puts some frequencies out of phase with other frequencies, resulting in phasing artifacts. We always see the frequency-response graph of equalizers, but rarely see the frequency versus phase graph, which is an inseparable part of an equalizer's operation. Figure 15.25 shows such a graph for a parametric equalizer.

Track 15.64: Kick *Q* No EQ
The source drums with no EQ on the kick.
In the next two tracks, the aim is to accent the kick's attack. Clearly, the two tracks do not sound the same, but they both achieve the same aim.

Track 15.65: Kick High Gain Narrow *Q*
The settings on the EQ are 3.3 kHz, +15 dB of gain, and a *Q* of 9. The resonant click caused by these settings might be considered unflattering.

Track 15.66: Kick Lower Gain Wider *Q*
The settings are 3.3 kHz, +9 dB of gain, and a *Q* of 1.3. The attack is accented here, yet in a more natural way than the previous track.

Plugin: Sonnox Oxford EQ
Drums: Toontrack EZdrummer

There are two important things we can learn from Figure 15.25. First, that the higher the gain is, the stronger is the phase shift. Second, we can discern that frequencies near the center frequency experience the strongest phase shifts. This behavior, although demonstrated on a parametric EQ, can be generalized for all other types of EQ—the more gain there is, the more severe phase artifacts become, with the greatest effect taking place around the set frequency.

The more gain, the more severe phase artifacts become.

One interesting question is: What happens with phase when we cut rather than boost? Figure 15.26 shows the same settings as in Figure 15.25, only this time for cut rather
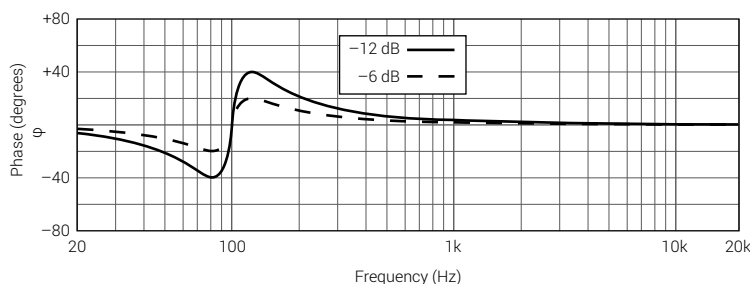
Figure 15.26 The phase response of a cut on a parametric equalizer. This is the same equalizer as in Figure 15.25, only this time with gain cut instead of boost.

than boost. We can see that the only difference is that the two response graphs are mirrored, but the phase extent remains the same. It is a myth that equalizers cause more phase shift when we boost—there is nothing in the science of either analog or digital filters to support such a claim. However, it is correct that we notice the phase shift more when we boost, for the simple reason that we also make the phase artifacts louder. It is therefore not surprising that many mixing engineers prefer to attenuate rather than boost when equalizing, and that many sources give such a recommendation. Also, when boosting we risk clipping.

When possible, it is better to attenuate rather than boost.

There is one type of digital equalizer that has a flat phase response—the **linear phase equalizer** (Figure 15.27). Digital filters are based on mathematical formulae. These formulae have stages, and the audio travels through the different stages. By making the formula of a filter symmetrical, as audio travels through one side of the formula its phase is shifted, but once it has gone through the mirrored side it shifts back in phase to its original inphase position. One issue with linear phase equalizers is that they require extensive processing power and a large buffer of audio (bigger than the typical 1,024 samples often provided by the audio sequencer). Thus, they are CPU-consuming and introduce plugin delay. Designers have to compromise between phase accuracy, processing power, and the delay the plugin introduces.

The phase artifacts produced by an equalizer are sometimes perceived by a trained ear as a subtle impairment of sound rather than as something very obvious. We do, however, often associate these artifacts with the term resonance.

Track 15.67: Snare Source
The source snare used for the following samples.
The following tracks all involve varying degrees of gain boost at 500 Hz. The EQ artifacts, which become more severe with gain, can be heard as a resonance of around 500 Hz.

Track 15.68: Snare 500 Hz 10 dB Up

Track 15.69: Snare 500 Hz 15 dB Up

Track 15.70: Snare 500 Hz 20 dB Up

Similar artifacts can be demonstrated using varying slopes of an HPF (represented in the following track names as dB/oct). Note that the applied filter is not of a resonant nature, yet a resonance around 500 Hz can be discerned:

Track 15.71: Snare HPF 500 Hz Slope 12

Track 15.72: Snare HPF 500 Hz Slope 24

Track 15.73: Snare HPF 500 Hz Slope 48

Plugin: Logic Channel EQ
Snare: Toontrack EZdrummer
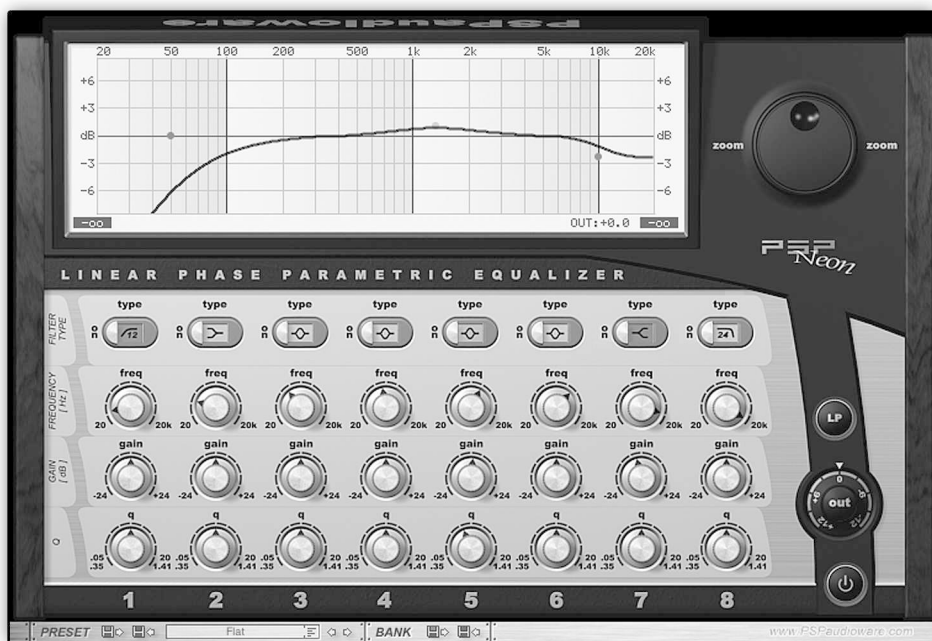
Figure 15.27 A linear-phase equalizer (the PSP *Neon*). This plugin offers eight bands with selectable filter type per band. It is worth noting the LP (linear phase) button, which enables toggling between linear phase and standard mode.

While linear-phase equalizers rectify one artifact of equalization, they do not rectify them all. Like standard equalizers, linear-phase ones have many other unwanted by-products such as ringing, lobes, and ripples.

Linear-phase equalizers may sound more "expensive." They are said to excel at retaining detail, depth, and focus. They are generally less harsh and likely to be more transparent. But many of the unwanted artifacts we encounter with standard equalizers (including those we associate with phase) are also a product of linear-phase equalizers. In addition, transients might not be well handled by the linear-phase process. In fact, there are situations where standard equalizers clearly produce better results. Thus, linear-phase equalizers provide an alternative, not a superior replacement.

Linear-phase equalizers are better at retaining detail, depth, and focus, but at the cost of processing power. Like standard equalizers, they can also produce artifacts, and thus provide an alternative, not a superior choice.

Track 15.74: Snare 500 Hz 20 dB Up
This track involves a linear phase equalizer with the same settings as in Track 15.70. The very similar resonance to that in Track 15.70 now has a longer sustain. In fact, its length is stretched to both sides of the hit. Many would consider the artifact on this track to be worse than the standard equalizer version on Track 15.70. Indeed, due to their design, when linear-phase equalizers produce artifacts, these tend to sound longer and more noticeable. This is a characteristic of all linear-phase equalizers, not just the one used in this sample.

Plugin: Logic Linear Phase Equalizer
Snare: Toontrack EZdrummer

The following tracks involve a comparison between linear-phase and standard EQ processing, and both are the result of boosting 12 dB on a high-shelving filter with its center frequency set to 2 kHz. Note how the highs on the standard version contain some dirt, while these appear cleaner and more defined on the linear-phase version.

Track 15.75: Guitar Standard EQ

Track 15.76: Guitar Linear Phase EQ

Plugin: PSP Neon HR

The frequency yin-yang

Figure 15.28 shows what I call the frequency yin-yang. I challenge the reader to identify the type of filter. Is it a +12 dB high-shelving filter brought down by 6 dB? Or is it a –12 dB low-shelving filter brought up by 6 dB? The answer is that it can be both, and, following our discussion about phase shift not varying between cut and boost, the two options should sound identical.
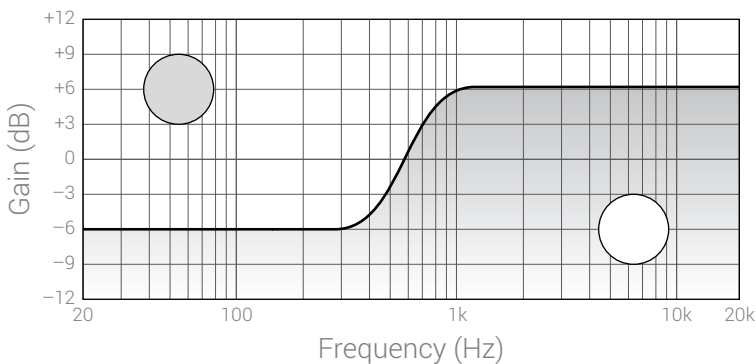
Figure 15.28 The frequency yin-yang.

Regardless of which of the two ways the frequency yin-yang is achieved, it teaches us something extremely important about our frequency perception: provided that the final signal is at the same level, boosting the highs or reducing the lows has the same effect. To make something brighter, we can either boost the highs or attenuate the lows. To make something warmer, we can boost the low-mids or attenuate from the high-mids up. While this concept is easily demonstrated with shelving filters, it works with other filters all the same. For example, to brighten vocals, we often apply HPF, and, despite removing a predominant low-frequency energy, the vocals can easily stand out more. Adding both highs and lows can be achieved by attenuating the mids. Sure, in many cases we have a very specific range of frequencies we want to treat, but it is worth remembering that sometimes there is more than one route to the same destination.

The following tracks demonstrate the frequency yin-yang shown in Figure 15.28*; the two tracks are perceptually identical:*

Track 15.77: Yin
This is the outcome of a high-shelving filter with a center frequency at 600 Hz, 12 dB boost, and –6 dB output level.

Track 15.78: Yang
This is the outcome of a low-shelving filter with a center frequency at 600 Hz, 12 dB attenuation, and +6 dB output level.

Track 15.79: Vocal Brighter
This is an equalized version of Track 15.15. What can be perceived as brightening is the outcome of an HPF (–6 dB point at 200 Hz, 12 dB/oct slope) and +1.8 output gain.

Track 15.80: Vocal Warmer
This equalized version of the previous track can be perceived as warmer. In practice, –4 dB was pulled around 3.5 kHz (and +1.3 of output gain).

Plugin: PSP Neon

One more example is necessary here. If we stop thinking about individual treatment, we can apply the frequency yin-yang at mix level as well. Say we want to add attack to a kick. Instead of boosting the attack frequency on the kick, we can attenuate the same frequency on masking instruments. In fact, a common technique based on this concept is called *mirrored equalization*. It involves a boost response on one instrument and mirrored response on another. This helps reduce frequency overlapping, and the same effect can be achieved with less gain on each filter. Figure 15.29 illustrates this.

Equalization and levels

By altering the level of a specific frequency range, we also alter the overall level of the signal. As per our axiom, louder-perceived-better, A/B comparisons can be misleading—we might think it is better when boosting on the equalizer and worse

Instrument 1
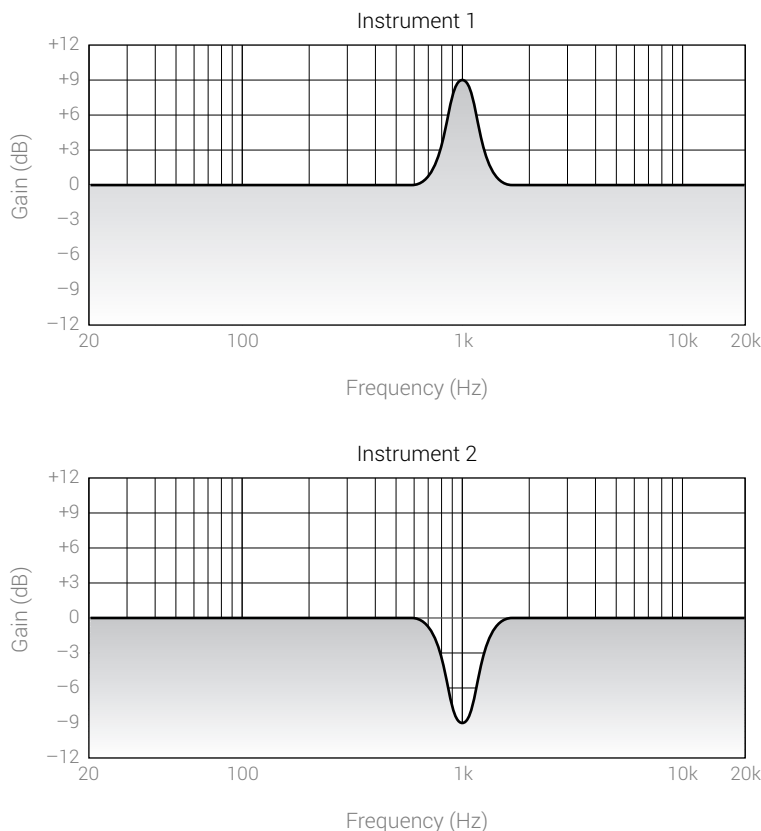
Instrument 2

Figure 15.29 Mirrored equalization. A boost on one instrument is reinforced
by a mirrored response on a masking instrument.

when attenuating. In order to allow fair A/B comparisons, some equalizers provide an output level control. After equalizing, we match the level of the processed signal to the unprocessed signal, so our critical judgment is purely based on tonality changes, not level changes. While it is understood why output level control is not found on the small EQ sections of even the largest consoles, it is unclear why many plugin developers overlook this important feature.

*Equalization alters the overall level of the signal and can deceive us, so we
think that boosts sound better.*

A/B comparison aside, the louder-perceived-better principle can lead to immediate misjudgments as we equalize, before compensating for the gain change. The risk is the same—we might think boosting improves the sound purely due to the overall level increase. The frequency yin-yang can minimize that risk; by taking the attenuation route, we are less likely to base our evaluation on the overall level factor. It was just mentioned that attenuating the lows of vocals could increase their presence. By attenuating a specific frequency range, we reduce masking on that specific range. Reducing the lows of the vocals, for example, would increase the low-end definition of other instruments. If by any chance the loss of overall level bothers us, we can always compensate for it. By boosting a specific frequency range, we increase

masking. The equalized instrument becomes more defined, but on the specific boosted range it masks more. This is why it is recommended to consider attenuation before boosting.

The psychoacoustic effect of taking away

In the first instance, our ears tend to dislike the attenuation or removal of frequencies. By attenuating or removing, we take away something and make instruments smaller. Our brains might need some time to get used to such changes. By way of analogy, it is like having a drastic haircut—it might feel a bit weird on the first day, but we get used to it after a while. Both listening in mix-perspective and giving the equalization effect some time to settle in can be beneficial when we take away some frequency content.

Attenuating or filtering frequencies might be right for the mix, but might not appear so at first.

Track 15.81: dGtr No EQ
The source track for the following samples, with no EQ applied.

Track 15.82: dGtr First EQ
The EQ settings in this track involve an HPF (144 Hz, 12 dB/oct) and a high-shelving filter (9 kHz, –8 dB). When played straight after the previous track, the guitar might appear smaller, but then playing the previous track again would reveal that the EQ in this track removed both rumble and high-frequency noise that existed on the previous track.

Track 15.83: dGtr Second EQ
This is an equalized version of the previous track, with the same EQ settings as in the previous track (an equalizer with the same setting was inserted serially). Again, this track appears smaller than the previous one. Also, comparing this track to the previous reveals that some high-frequency noise still existed in the previous track.

Track 15.84: dGtr Third EQ
This is an equalized version of the previous track, this time involving a band-pass filter (between 100 Hz and 2 kHz). Now Track 15.82 sounds bigger compared to this one. Comparing this and the unequalized track would make the latter sound huge.

Plugin: Digidesign DigiRack EQ 3

One specific technique that can be used to combat this psychoacoustic effect is making an instrument intentionally smaller than appropriate, forgetting about the equalization for a while, then going back to the equalizer and making the instrument bigger. This could allow a fairer judgment when deciding how to equalize as we make things bigger rather than smaller.

Applications of the various shapes

Having the choice between pass, shelving, and parametric filters, we employ each for a specific set of applications. The most distinct difference between the various shapes puts pass and shelving against the parametric filter. Both pass and shelving filters

affect the extremes of the frequency spectrum. Parametric filters affect a limited, often relatively narrow bandwidth, and rarely do we find them around the extremes. In a more specific context, the basic wisdom is this:

- **Pass filters**—used when we want to **remove** content from the extremes; for example, low-frequency rumble.
- **Shelving filters**—used when we want to alter the overall **tonality** of the signal (partly like we do on a hi-fi system), or to **emphasize** or **soften the extremes**; for example, softening exaggerated low-frequency thud.
- **Parametric filters**—used when what we have in mind is a **specific frequency range** or a **specific spectral component**; for example, the body of a snare.

It is sensible to introduce the three different types into a mix in this order: first use pass filters to remove unwanted content, then use shelving for general tonality alterations, and finally use parametric filters for more specific treatment. The following sections detail the usage of each filter type.

HPFs

HPFs are common in mixes of recorded music. First, they remove any low-frequency gremlins such as rumble or mains hum. Then, recorded tracks can contain a greater degree of lows than needed in most mixes (sometimes due to the proximity effect; sometimes due to the room). When the various instruments are mixed, the accumulating mass of low-end energy results in muddiness, lack of clarity, and ill definition. HPFs tidy up the low end by removing any dispensable lows or low-mids. Doing so can clear some space for the bass and kick, but more importantly it can add clarity and definition to the treated instrument. This is worth stressing:

Despite removing spectral content, HPFs increase clarity and definition and can make the treated instrument stand out more in the mix.

An HPF might be applied on every single instrument. Vocals, guitars, and cymbals are common targets. Vocals, nearly by convention, are high-passed at 80 Hz—a frequency below which no contributing vocal content exists. Higher cut-off frequencies are used to remove by-products of the proximity effect or some body that might not have a place in the mix. Many guitars, especially acoustics, occupy unnecessary space on the lows and low-mids, which many other instruments can use instead. When acoustic guitars play a rhythmic role, they are often filtered quite liberally with most or all of their body removed. Cymbal recordings often involve some drum bleed that cannot be gated (e.g., removing the snare from underneath the ride), so the filter also acts as a spill eliminator. Pianos, keyboards, snares, and any other instrument can benefit from low-end filtering all the same. We sometimes even filter the kick and bass in order to mark their lowest frequency boundary, and in turn that of the overall mix.

The frequencies involved in this tidying up process are not strictly limited to the lows. Cymbals, for example, might be well within the low-mids. A possible approach is to simply sweep up the cut-off frequency until the timbre of the treated instrument breaks, then back it off a little. Usually, the busier the mix, the higher frequency HPFs reach. Over-filtering can create a hole in the frequency spectrum or reduce warmth. To rectify this, we can pull back the cut-off frequency on some instruments.

One interesting characteristic of HPFs is that they can be pushed quite high before we notice serious impairment to the treated instrument. The reason for this has to do with our brains' ability to reconstruct missing fundamentals. What an HPF removes, the brain reconstructs—we clear space in the mix, but do not lose valuable information.

Due to the brain's ability to reconstruct missing fundamentals, HPFs can be used quite generously.

Track 15.85: No HPF (aGtr)
The source track for the following samples.

In the following tracks, a 12 dB/oct HPF is applied with different cut-off frequencies (denoted in the track names). Virtually all of the following degrees of filtering might be appropriate in a mix. Note how, despite filtering the fundamentals of the notes, our ears have no problem recognizing the chords.

Track 15.86: HPF 150 Hz (aGtr)

Track 15.87: HPF 250 Hz (aGtr)

Track 15.88: HPF 500 Hz (aGtr)

Track 15.89: HPF 1 kHz (aGtr)

Track 15.90: HPF 3 kHz (aGtr)

Plugin: McDSP FilterBank F2

HPFs can manipulate the characteristics of reverbs. Generally speaking, the size, depth, and impression of a reverb all focus on the lower half of the spectrum (lows and low-mids). The highs mostly contribute some definition and spark. By filtering the lows, we can reduce the size of a reverb and its resultant depth. The higher the cut-off frequency is set, the smaller the size becomes. Although we have full control over the size and depth of a reverb when using a reverb emulator, these factors are imprinted into a reverb captured on recordings.

An HPF can reduce the size and depth of reverbs.

In the following tracks, percussion is sent to a reverb, which is filtered by an HPF with various cut-off frequencies (denoted in the track name). Note how the dimension of the space shrinks the higher the cut-off frequency is (this is also the consequence of the overall level attenuation caused by the filter).

Track 15.91: No HPF (Reverb)

Track 15.92: HPF 400 Hz (Reverb)

Track 15.93: HPF 1 kHz (Reverb)

Track 15.94: HPF 2 kHz (Reverb)

Track 15.95: HPF 4 kHz (Reverb)

Track 15.96: Percussion Dry
This is the dry percussion used in the previous samples. Compare this track to the previous one and note the appealing enhancement the filtered reverb provides.

Plugin: PSP MasterQ, Audio Ease Altiverb
Percussion: Toontrack EZdrummer

HPFs are also used to remove pops—low-frequency air bursts caused by Ps, Bs, or any other plosive sounds such as an emphasized "who." The higher the cut-off frequency is, the more the pop is reduced, but also the more likely the timbre is to suffer. In the days of tapes, removing these pops required filtering the whole vocal track and compromises had to be made. Audio sequencers offer a much more elegant solution, filtering the pop only. This is achieved using offline processing, where we simply select each pop, apply filtering, and then crossfade to prevent clicks. Figure 15.30 illustrates this.

Figure 15.30 Pop removal using offline processing. The top track, before processing, shows an evident pop. An offline HPF is set with a cut-off frequency of 368 Hz. Had the filter been applied in real time throughout the track, much of the vocal's thickness would have been lost. The offline process was only applied on a limited selection that included the pop and some margins for the crossfades. The crossfades were then applied to prevent clicks. The resultant bottom track shows the apparent energy loss of the pop.

Track 15.97: Vocal Pop Raw
The raw vocal recording, involving a pop on "**P**ath."

Track 15.98: Vocal Pop Removed
The vocal pop is removed following the offline processing shown in Figure 15.30.

Track 15.99: Vocal Pop EQ Throughout
This track is the result of applying exactly the same settings throughout the vocal take using a real-time plugin. Not only is the body of the vocal reduced, but also the pop becomes evident again.

Plugin: Digidesign DigiRack EQ 3

LPFs

Unlike their lofty twin, LPFs are somewhat neglected when it comes to mixing. Unfortunately, our ears do not reconstruct the higher harmonics once lost, so we find the removal of high frequencies more disturbing. Some instruments get away with it. For example, filtering above 10 kHz from a kick could be unnoticeable, despite the fact that some energy does exist there. Instruments with richer harmonic or overtone content are more sensitive to low-pass filtering.

LPFs are used for two principal tasks. The first is the removal of **hiss** or **high-frequency noise**. The second is marking the high-frequency boundary of a specific instrument. An HPF and LPF are occasionally combined to form a flexible band-pass filter (Figure 15.31).

The general idea is to restrict an instrument to a limited band, which can increase separation. One example of where this can be beneficial involves distorted guitars. These often have an excess of both low-end rumble and high-end noise. Being forceful maskers as they are, band-limiting them also reduces their masking inter-action to a limited range of frequencies. Other instruments can also benefit from such an approach.

Figure 15.31 An HPF and LPF combined to form a band-pass response (the MOTU *MasterWorks EQ*). The two filters are set to remove frequencies below 200 Hz (HPF) and above 5 kHz (LPF). Band-limiting instruments in a similar way can increase separation and reduce masking.

Shelving filters

While both pass and shelving filters affect the extremes, pass filters remove, while shelving filters soften or emphasize. Needless to say, a shelving filter can boost and a pass filter cannot. One important characteristic of shelving filters is that they are not as aggressive as pass filters—while with shelving filters we have ultimate control over the gain amount, with pass filters we get the less flexible slope control. Shelving filters can be set to have very little effect and their response curves are generally gentler.

Shelving filters can provide a softer alternative to pass filters.

The various applications of pass filters can easily become a shelving job if we exchange the word "remove" for "soften." For example, sometimes we only want to soften the body of vocals, not get rid of it altogether. Sometimes we want to reduce the size of a reverb, but still keep some of its warmth. The high-frequency noise on distorted guitars might only call for moderating. Sure enough, we can always use a combination of pass and shelving—pass to remove unwanted content (e.g., rumble), shelving to soften wanted content (e.g., body).

Shelving filters are often associated with the terms *thick* and *thin* (for the lows), and *bright* and *dark* (for the highs). We are all familiar with the shelving effect from our hi-fis. Having the same ability to shape the tonality of each individual instrument is a great asset. Two key challenges that the frequency domain involves are balanced spectrum and separation between various instruments. Referring back to Figure 6.2, it was suggested that we take an abstract approach and imagine where each instrument is located on the frequency spectrum, how wide it is, and whether there are any lacking frequency ranges in our mix. In the journey to balance spectrum and separation, we might want to nudge various instruments up or down on the frequency spectrum. We might also want to narrow or widen their frequency range. This practice is often referred to as **tuning** an instrument into the frequency spectrum. Keeping the same abstract approach, we could really use the ability to nudge the lower or higher boundary of an instrument.

Shelving filters are great tools for doing this, and Figure 15.32 illustrates how. Essentially, if you want an instrument to take more of the highs, boost using high shelving; if you want it to take less, attenuate. Using the very same principle, if you want an instrument to
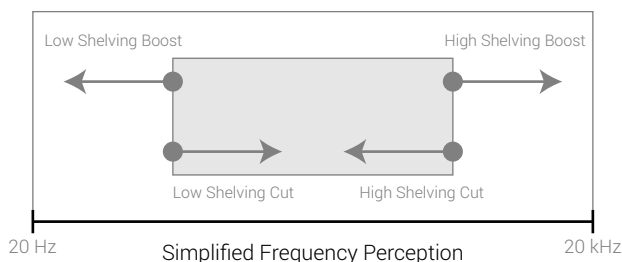
Figure 15.32 Tuning an instrument frequency range using shelving filters. Low shelving affects the low boundary: boost lowers it, cut makes it higher. High shelving affects the high boundary: boost makes it higher, cut lowers it.

take more of the lows, boost on a low shelving; for less of the lows, attenuate. The only trick is to hit the right shelving frequency. The advantage that shelving filters have over pass or parametric ones is that they can be set for a very subtle effect (unlike pass) and affect a wide frequency range with a flat response (while parametric are more oriented toward a specific range of frequencies).

Shelving filters let us tune instruments into the frequency spectrum, help create separation, and balance the overall frequency spectrum.

Track 15.100: Drums No EQ
The source track for the following samples, in which only the snare is equalized.

Track 15.101: Drums HSF Up
A 3 dB boost at 1.8 kHz on a high-shelving filter creates the impression that the snare's high boundary has been extended up in the frequency spectrum.

Track 15.102: Drums HSF Down
A –4 dB at 1.8 kHz on a high-shelving filter creates the impression that the snare's high boundary has moved down in the frequency spectrum.

Track 15.103: Drums LSF Up
A –3 dB at 170 Hz on a low-shelving filter creates the impression that the snare's low-frequency boundary has moved higher.

Track 15.104: Drums LSF Down
A 4 dB boost at 200 Hz creates the impression that the snare's low boundary has moved lower.

Plugin: Digidesign DigiRack EQ 3
Drums: Toontrack EZdrummer

Parametric filters

Parametric filters are used when we wish to treat a specific range of frequencies. Often these specific ranges are associated with some spectral component of the treated instrument (e.g., body) or with a general sonic quality (e.g., nasal). Although not limited to the midrange, often this is the range parametric filters treat.

Recognizing and comprehending the essence of various frequency ranges takes a lifetime of experience. A popular trick can help us locate the right **frequency** when using parametric filters. Say we are trying to find the attack of a kick. The technique involves a generous boost with a narrow *Q*, and then sweeping the center frequency across the spectrum. The frequency at which the kick's attack is most profound is the frequency we are after. Then, we adjust the gain and *Q*. While many mixing resources mention this technique, it has some strong opposition as well. For one, it hinders critical evaluation skills, which are based on the process of listening, analyzing, and then taking action. It can be compared with the less-experienced guitarist playing all sorts of possible chords in the search for the right one. Then, there is the same risk of

liking the boost too much, and ending up with a higher gain than appropriate. Having said that, this technique can still be helpful for the less-skilled engineer and with tricky material.

Although **gain** settings are often straightforward, even here there is more than meets the eye. We have already discussed the upward versus sideways principle, where we demonstrated that gain can sometimes be traded for Q. How much effect an equalizer has is a function of both the gain and Q settings. Perhaps surprisingly, exactly the same settings and response curve on two different equalizers might not produce identical results. While the gain on an equalizer can have a noticeable effect with 0.5 dB of boost, certain equalizers offer up to 18 dB of gain. A 15 dB boost could be harrowing, but it can also be musical. For these reasons, it is not easy to generalize what typical gain settings are—it depends on many factors, such as the frequency, Q, source material, and equalizer being used. Some people regard 1 dB as gentle, 3 dB as conservative, 6 dB as moderate, 9 dB as serious, and 12 dB as extreme.

One useful equalization technique involves distributing the equalization load across different frequencies. More specifically, instead of equalizing around one frequency only, we also equalize around its harmonic (an octave above). Figure 15.33 demonstrates this. This lets us reduce the gain on the first frequency and shift that gain to the harmonic. Compared to using one band only, each of the two bands will have a less drastic response and less phase artifacts. Also, having the equalization take place at two points often has the advantage of overcoming masking.

Harmonics equalization can reduce artifacts and can be less prone to masking.
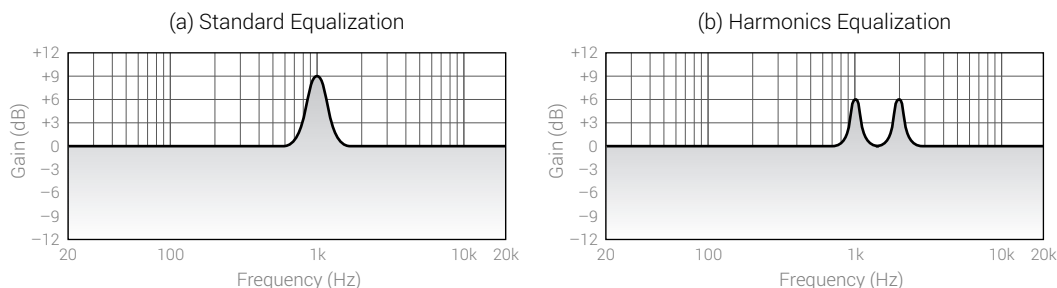
Figure 15.33 Harmonics equalization. (a) Standard equalization, where a single range of frequencies is equalized. (b) Harmonics equalization, where the same range is equalized with a softer response, but its harmonic is also equalized.

The source track for the following samples is Track 15.64. The aim is to accent the kick's attack.

Track 15.105: Kick Standard EQ
This track involves a single-band parametric filter set to 15 dB of boost around 2 kHz, with Q set to 6.7. This results in an evident click, but also a noticeable resonance around 2 kHz.

> **Track 15.106: Kick Harmonic EQ**
> Instead of equalizing a single frequency, three parametric filters are employed with their center frequency set to 2, 4, and 8 kHz. Each band is set to around 8 dB of boost, and all *Q* settings are 6.7. The result is still an accented click, but one that is cleaner and less obstructive than in the previous track.
>
> Plugin: Sonnox Oxford EQ
> Drums: Toontrack EZdrummer

Q settings are also a vital part of equalization. A *Q* too wide might take on board unrelated frequencies, whereas a narrow *Q* might not be effective enough. A narrow *Q* is often used when we want to treat a very narrow bandwidth, often an unwanted noise or a resonant frequency. A wide *Q* tends to sound more natural and is mostly used to emphasize wanted parts of the frequency spectrum and for fine tonal alterations. Accordingly, a wide *Q* is more common with boosts and a narrow *Q* with cuts.

Equalizing various instruments

Spectral components

We say that we can divide the timbre of each instrument into four parts: lows, low-mids, high-mids, and highs. Learning how each of these bands affects the tonality of an instrument is one of the first things mixing engineers should grasp. It only takes a few minutes to experiment with how each band affects the tonality of an instrument—the experience is invaluable. However, the equalization of each instrument is not as simple as that—subtle alteration of very specific frequencies can have a dramatic effect. By way of analogy, the equalization of each instrument is like a safe waiting to be cracked. It would be fair to make no subdivisions at all—it is simply 20 Hz–20 kHz and what we are after can be anywhere. Once familiar with the four different bands, we progress and learn how smaller ranges affect the timbre of various instruments. Each instrument has a few influential spectral components—frequency ranges highly vital for its timbre. Learning these takes a bit more time as different instruments have different spectral components at different ranges.

The rest of this chapter covers common instruments, their spectral components, and other frequency ranges that might be of interest to us.

The information presented in this chapter hereafter, including any frequencies and advice, constitutes guidelines only. Each recorded instrument produces different sounds; each microphone and its position capture different frequency content. The spectral complexity of the sounds we treat is immense. The true essence of equalization cannot be realized by mere guidelines—we must develop critical evaluation skills, which, in the case of equalization, takes time. I urge the reader to experiment with the information presented below, rather than take it as dogma.

Vocals

It can be argued that the term *vocals* alone cannot encompass the great variety of human voices. For one, the voices of males and females are different. Then, each person has so many unique qualities in their voice that every singer can be considered a completely different instrument. It should be fair to assume that vocal equalization starts by eliminating the wrongs—the likes of muddiness, honky sound, and sibilance if apparent (which deessers are likely to rectify better). Being the key element of many productions, vocals must be most prominent. It is vital for vocals to overcome any instrument that masks them. This might entail equalization not only on the vocals, but also on the masking instrument. Different vocal parts have different roles in mixes, so at times we make them sweet, at times warm, and at times aggressive. Generally speaking, the more powerful the vocal part is, the more equalization will be used.

Figure 15.34 shows the possible frequency range of the human voice. The long gradient bar shows the possible frequency content of both fundamentals and harmonics. The bounded rectangle denotes the possible fundamentals range, which ranges between 82 and 1,398 Hz (E2–F6). No person can produce this full range of notes. Also, the extremes of this range can only be produced by skilled opera singers. To give an example, Mariah Carey's vocal range spans approximately 123–830 Hz (B2–G#5). Typical pop singers cover a far less impressive range than that. Above and below the bar are frequency ranges relevant to vocals. It is worth noting that equalizing a specific range can yield either positive or negative results. For example, by boosting the high-mids we might add presence, but at the same time add unwanted sibilance. By attenuating the low-mids we might add clarity, but lose warmth and make the sound thin.
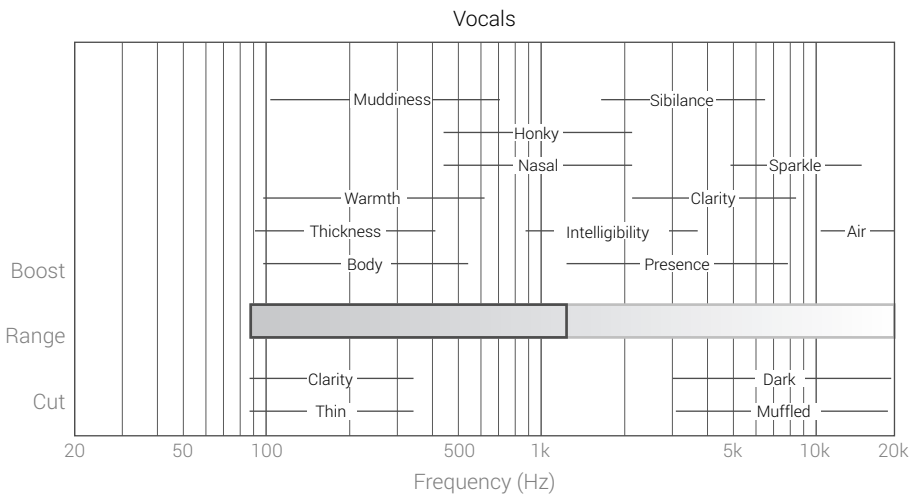
Figure 15.34 The frequency range of the human voice and relevant frequency ranges. The gradient bar denotes the range of possible fundamentals and harmonics. The possible fundamental range is confined in a rectangle. The associations above the bar are likely to be emphasized by a boost in their frequency ranges, the associations below by a cut.

Track 15.107: Vocal No EQ
The unequalized track.

Track 15.108: Vocal HPF 280 Hz
A 12 dB/oct HPF at 280 Hz removes some muddiness but also some body. The vocal here is clearer compared to the previous track.

Track 15.109: Vocal HSF 9 dB Up
This is an equalized version of the previous track. A 9 dB boost at 10 kHz adds some air and sparkle to the vocal.

Track 15.110: Vocal 1 kHz Boost
This is an equalized version of Track 15.108. An 11 dB boost at 1 kHz adds some nasal character to the voice.

Track 15.111: Vocal 1 kHz Boost
This is an equalized version of Track 15.108. A 6 dB boost at 3 kHz adds some presence and clarity.

Track 15.112: Vocal 300 Hz Dip
This is an equalized version of Track 15.108. A 9 dB dip at 300 Hz also adds some clarity.

Track 15.113: Vocal 130 Hz Boost
This is an equalized version of Track 15.108. An 8 dB boost at 130 Hz adds some warmth.

Plugin: Sonnox Oxford EQ

Overheads

The overheads glue all the individual drums together. When the overheads are too quiet, the drums can sound very artificial and come across as lifeless. There are three general approaches to the involvement of overheads in mixes. What differs from one approach to the other is the amount of low-end filtering in combination with the overheads' level:

- **Main stereo pair**—the overheads are responsible for the majority of the drum sound, with possible minor reinforcement from close-mics, mostly kick and snare. This is typical in conventional jazz mixes and more natural-sounding productions. No or subtle low-end filtering is applied.
- **First among equals**—an equal share between the overhead and the close-mics. This is common in contemporary productions where kicks and snares are very prominent. An HPF might be positioned in the frequency range between the kick and the snare, mostly to clear the low end for the kick's close-mic and reduce spatial cues that send the various drums backward.
- **Cymbals only**—the overheads are responsible for the cymbals only, and the close-mics for all other drums. This least-natural approach entails strong high-pass filtering (even into the high-mids), intended to remove as much drum content (and room) as possible. This technique can be a lifeboat when the overheads are seriously flawed—nothing new in low-budget and home recordings.

In addition to the drum kit itself, overhead recordings also include the room response (the reverb). This creates an illusion of space in the mix, which is more or less similar to
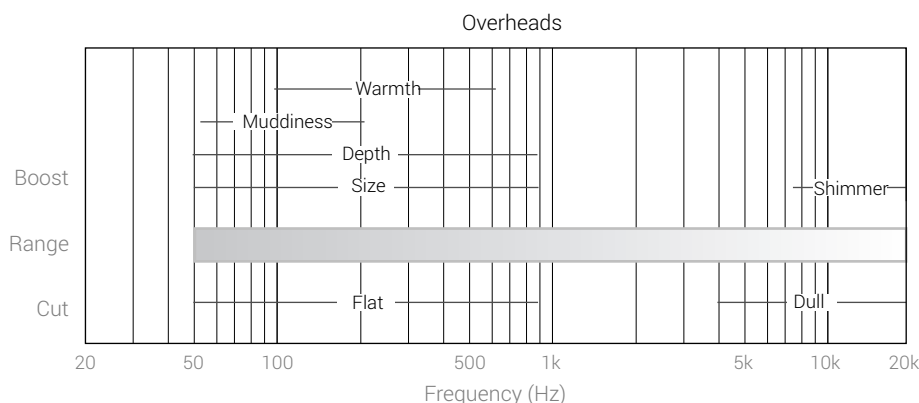
Figure 15.35 The frequency range of overheads and relevant frequency ranges.

that of the recording space. The problem is that sometimes there is too much room, sometimes the space is just not appealing—most domestic rooms and even some studio spaces can produce a weary mix ambiance. HPFs let us correct, to some extent, such recordings.

The size, depth, and impression of reverbs focus on their lows. One particular problem is that low-frequency drums, notably the kick, tend to excite most rooms more than, say, the cymbals. This is normal as low frequencies take longer to absorb. Yet, it is a definite attribute of overheads' recordings due to the small live rooms often used. As an outcome, the reverb from the overheads imposes some depth on the kick, toms, and the snare. Such depth might be unwanted, and, no matter how the close-mics are mixed, it can be hard to get the drums to the in-your-face position. Filtering the lows from overheads reduces spatial information and helps bring the individual drums forward in the depth field. The more we filter, the smaller we make the room and the closer the image becomes. In order to eliminate more of the individual drums but retain more of the cymbals' timbre, a steep filter is often needed. Combined with the fact that overheads are a broadband signal that involves transients, a high-quality HPF is a requisite for the task. For more gentle results, a shelving filter can be used instead of a pass filter—it allows us to keep more of the warmth and spatial impression that the overheads contribute, but still reduces their size and depth contribution.

By taking the cymbals-only approach, we can filter a flawed room from the overheads. If this dries out the mix ambiance, we can always use a technique called **ambiance reconstruction**—sending the individual drums into a reverb emulator to re-create the missing ambiance. We have full control over the reverb we choose, but we have to make sure that it matches the high-frequency reverb still present on the overheads.

Unless recorded in a good room, the low end of the overheads might muddy the mix. Both the lows and low-mids can benefit from some attenuation that will clear some space for the kick (close-mic), bass, and the fundamentals of other instruments. However, caution must be taken as the same lows might contribute to general warmth and the sense of ambiance in the mix.

The highs of the overheads contribute some shimmer. If no cymbal mics are mixed, the overheads might be the ones to govern the high end of the mix. Thus, how bright or dull we make the overheads can also govern how bright or dull the mix is, and in turn how balanced the overall mix is.

Track 15.114: Overheads No EQ
The unequalized track, used for all the following samples.

Track 15.115: Overheads HPF 100 Hz
An 18 dB/oct with –9 dB point at 100 Hz. The filter removes much of the ambiance and attenuates the kick.

Track 15.116: Overheads HPF 420 Hz
A steep 36 dB/oct HPF is set to 420 Hz. This type of processing could be used for the cymbals-only approach as it filters much of the kick and the snare. Depending on the mix, even cut-off frequencies higher than 420 Hz might be appropriate.

Track 15.117: Overheads HSF Up
A 2.4 dB boost at 7.8 kHz adds shimmer.

Plugin: Universal Audio Cambridge Equalizer and Filters
Drums: Toontrack EZdrummer

Kick

If you listen to a collection of productions, you will find kicks in many shapes and forms—the thunderous, the rounded, the 909, the 808, the woody, the typewriter style, the basketball, the pillow-like, to name a mere few. How we shape the sound of a kick is often part of creative decisions we make; we control aspects such as solidness, punchiness, thickness, and snappiness. But the sound of the kick is also a practical affair. Being the admiral of the rhythm, the kick has immense weight in the ability of the mix to move people. A weak kick on a hip-hop mix means a weak hip-hop tune; ditto for dance; and mostly so for rock, pop, and metal. When it comes to mixing kicks, equalizers play a powerful role.

Kicks have two main spectral components—the impact and the attack. These two components correspond closely with the lows and high-mids, respectively. Within these two ranges, different frequencies have a different effect. For example, 60 Hz might produce more oomph, while 90 Hz might produce more thud. Dance and hip-hop might benefit from robust oomph, while rock and pop might benefit from a solid thud. On the high-mids, the higher the frequency, the more the click. A healthy boost around 8 kHz can produce the typewriter click that many associate with heavy metal. The low-mids of the kick could have little to offer, and attenuating them can clear useful space for the low-order harmonics of other instruments. On the highs, kicks have very little valuable energy, which is sometimes entangled with hiss. Rolling off the highs from a kick could go unnoticed.
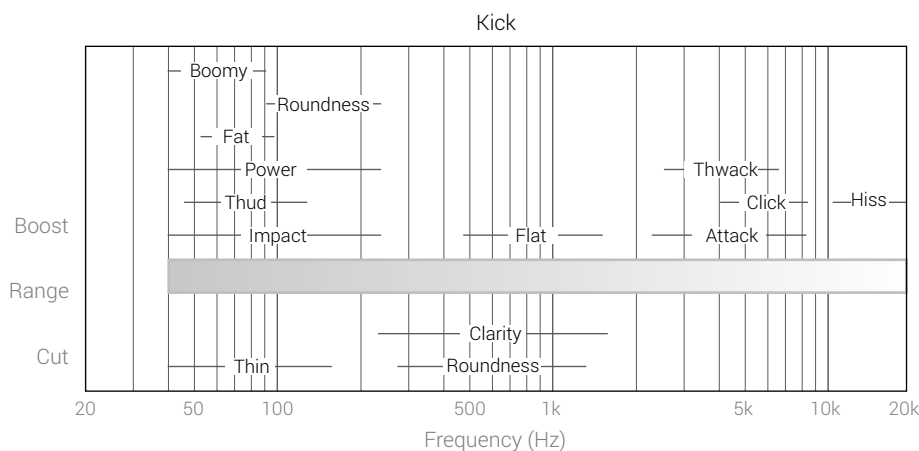
Kick

Figure 15.36 The frequency range of a kick and relevant frequency ranges.

One important aspect of equalizing a kick is how it interacts with the bass. As the two instruments are competing for the lows, sometimes we have to sacrifice the impact of one for the benefit of the other. Another important issue with mixing kicks is that there is always a risk of mishandling the low end due to limited monitors or poor acoustic environments. Under such conditions, it pays to spend some time going through the stabilizing options described in Chapter 4.

Track 15.118: Kick No EQ
The unequalized track, used for all the following samples.

Track 15.119: Kick LSF 160 Hz Up
A 6 dB boost at 160 Hz on a low-shelving filter adds some thud.

Track 15.120: Kick 344 Hz Dip
This is the result of –10 dB at 344 Hz.

Track 15.121: Kick 1600 kHz Boost
The added thwack on this track was achieved by a 9 dB boost at 1.6 kHz.

Track 15.122: Kick 5 kHz Boost
A 6 dB boost at 5 kHz adds some attack.

Track 15.123: Kick 7 kHz Boost
A 6 dB boost at 7 kHz also adds some attack, but with more click.

Track 15.124: Kick Two Bands
This track is the outcome of –8 dB at 270 Hz on a high-shelving filter, together with –10 dB at 261 Hz.

Plugin: Universal Audio EX-1
Kick: Toontrack EZdrummer

Snare

Perhaps the prince of rhythm, the tonality of the snare is also an important aspect of many mixes. Snare recordings are not always ideal. It is often only high-budget projects that can afford snare comparisons, fresh skins, accurate tuning, appropriate muffling, a suitable microphone, and correct placement. In all other projects, snares are prone to a few issues that might need correction—excess rattle and resonant ring are two such examples.

Although the spectral components of a snare are not as distinct as those of a kick, we often talk about body and presence, while also considering aspects such as snap and crispiness. A part of snare mixing entails tuning them into the frequency spectrum; that is, deciding how high or low they sit, also with relation to other instruments. A dark snare will be more distant and laid back. A bright snare will be more upfront. Automating snare sounds (between verse and chorus, for example) has rationale in more than a few mixes.
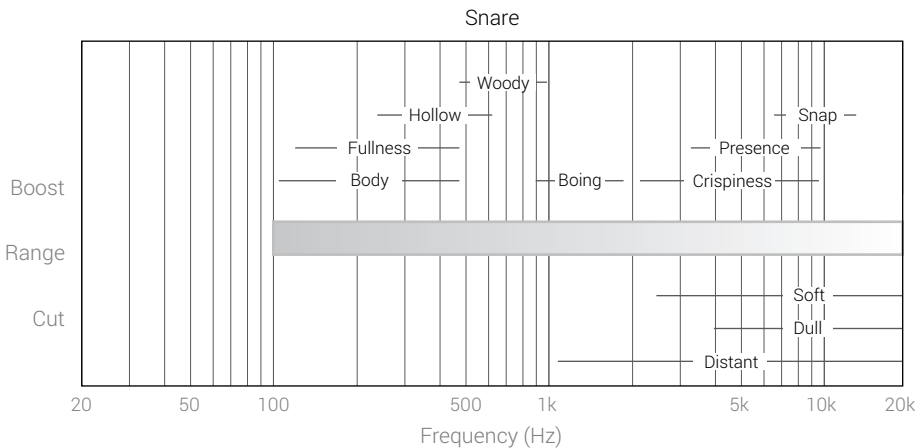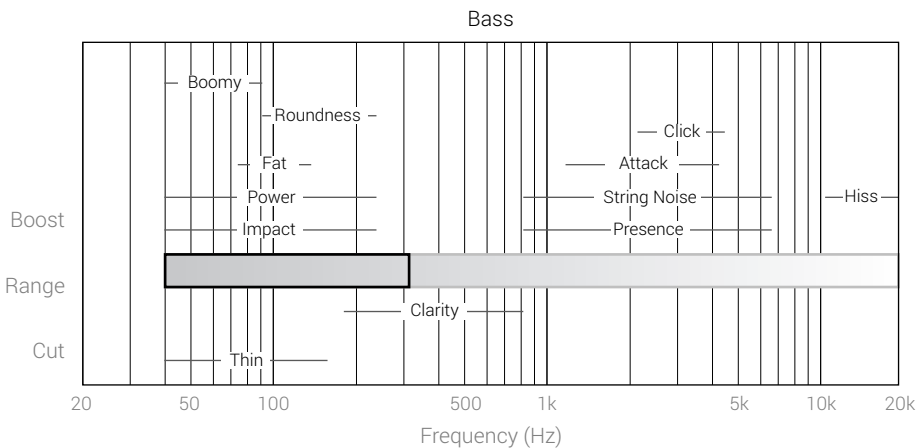

Figure 15.37 The frequency range of a snare and relevant frequency ranges.

Track 15.125: Snare No EQ
The unequalized track, used for all the following samples.

Track 15.126: Snare 150 Hz Boost
A 6 dB boost at 150 Hz adds body.

Track 15.127: Snare 400 Hz Dip
This is the result of a –6 dB at 400 Hz.

Track 15.128: Snare 910 Hz Boost
A 6 dB boost at 910 Hz.

Track 15.129: Snare 4700 Hz Dip
A –6 dB at 4.7 kHz.

Track 15.130: Snare HSF 10 kHz Boost
A 4 dB boost at 10 kHz on a high-shelving filter.

Plugin: PSP Neon
Snare: Toontrack EZdrummer

Toms

While often only playing occasionally, toms might just be blended into the overall drum sound, but can also be mixed to have a distinct powerful impact. Toms, notably floor toms, can get a bit out of control unless we contain their lows. Also, it is important to observe timbre differences between the various toms—poorly tuned toms that are recorded with different microphone brands can sound as if each belongs to a different kit.

Toms have some similarity to kicks in the way we mix them, but they have more defined pitch and longer decay. Unsurprisingly, a high tom will have its spectral components higher on the spectrum than a floor tom. At the very lows, there is usually rumble that might need filtering. The fullness and thud are around the higher lows, 200 Hz or so. The attack, as with kicks, is on the high-mids. It is also worth seeing whether cutting the highs reduces cymbal spill without affecting the sound of the toms.
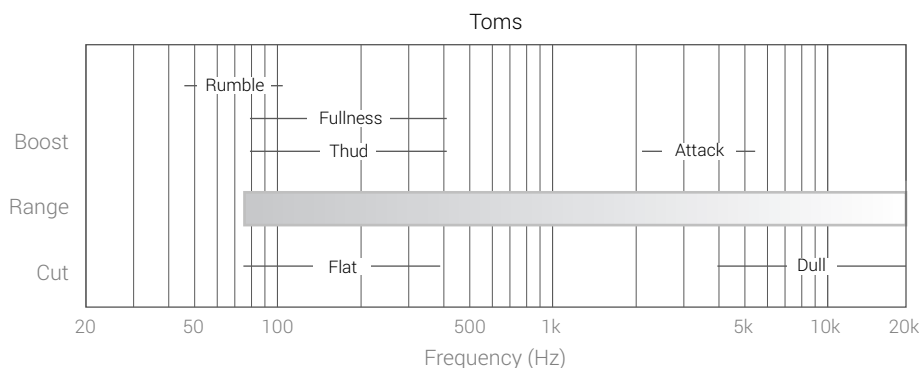
Figure 15.38 The frequency range of toms and relevant frequency.

Track 15.131: Tom No EQ
The unequalized track, used for all the following samples.

Track 15.132: Tom 100 Hz Boost
A 9 dB boost at 100 Hz adds some thud and tone.

Track 15.133: Tom 400 Hz Dip
A –7 dB at 400 Hz.

Track 15.134: Tom 3 kHz Dip
A –9 dB at 3 kHz reduces the attack.

Plugin: McDSP G Equalizer
Tom: Toontrack EZdrummer

Hi-hats and cymbals

The hi-hats are often the brightest element constantly playing in the mix. As such, they could play a big part in our perception of how bright the mix is altogether. One common problem with hi-hats is that they sound detached from the rest of the mix. To elaborate, they are perceived in the spectrum as the brightest instrument, but there is some empty frequency range below them. This is often the outcome of overemphasized highs. One characteristic of many pleasing mixes is that the hats are not glaring bright—they just shimmer. It is worth remembering that it might be the overheads contributing to this aspect in the mix and not close-mics.
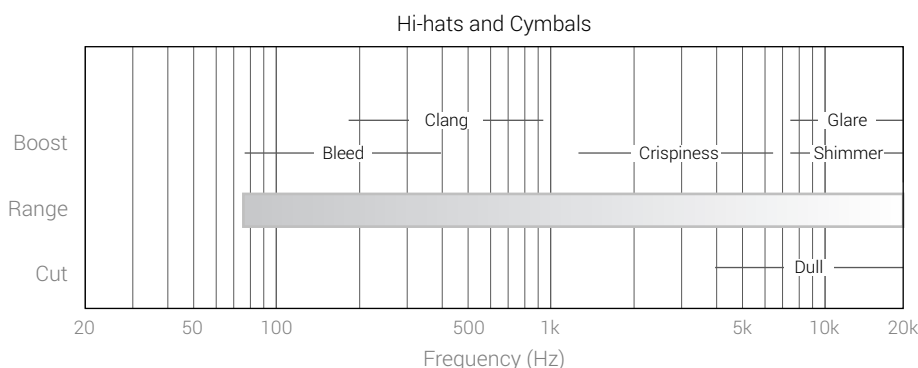
Figure 15.39 The frequency range of hi-hats and cymbals, and relevant frequency ranges.

Cymbals have very similar characteristics to hi-hats, but as they only play occasionally they play a less crucial part in the overall frequency balance. The larger the cymbal, the lower is its fundamental. Hi-hats, for example, can have some content below 500 Hz; rides go lower than that. These lower ranges of cymbals, which often involve bleed, can sometimes be pulled without affecting too much of the timbre.

Track 15.135: Cymbals No EQ
The unequalized track, used for all the following samples.

Track 15.136: Cymbals HPF 400 Hz
An 18 dB/oct HPF at 400 Hz. Such a treatment would doubtfully be noticeable with the rest of the mix playing along.

Track 15.137: Cymbals 400 Hz Boost
A 6 dB boost at 400 Hz adds some clang, mainly to the ride.

Track 15.138: Cymbals 1 kHz Boost
A 4 dB boost at 1 kHz adds some crisp.

Track 15.139: Cymbals HSF 6 kHz Up
A 4 dB boost at 6 kHz on a high-shelving filter adds some glare.

Plugin: Digidesign DigiRack EQ 3
Cymbals: Toontrack EZdrummer

Bass

Usually one of the trickiest instruments to mix is the bass—whether a guitar or a synth. For one, recordings tend to be very different from one another. But the real problem with bass sounds is that their tone often changes in relation to the note being played. For example, the lower E (41 Hz) produces substantially more low frequencies than the E two octaves above it (164 Hz). Solid tone can be achieved with equalizers, although both dynamic equalizers and compressors (especially multiband) can do a better job. There are also different playing techniques and different bass sounds in different genres. The bass in funk, for example, is miles away from the bass in metal.

There are two principal aims when mixing a bass: power from solid lows, and definition. Neither is easy to achieve. Being of a sustained nature in most productions, the bass is often the instrument that fills the low end. Too many lows and the mix is boomy; too few and the mix is thin (and, as mentioned, this characteristic can change in relation to the note being played). Then, the bass and the kick often compete for the lows. Most bass sounds have relatively little energy on the high-mids and highs, and some have no energy at all above, say, 5 kHz (specific recordings can have almost nothing even below that). On the low-mids, there are usually enough instruments fighting for space. So defining the bass might require some experimentation. Any mixing tool that generates harmonics (distortions and enhancers, for instance) can be extremely useful in providing some definition for bass sounds.

Figure 15.40 The frequency range of a bass guitar and relevant frequency ranges.

Track 15.140: Bass No EQ
The unequalized track, used for all the following samples.

Track 15.141: Bass LSF 230 Hz Boost
A 7 dB boost at 230 Hz on a low-shelving filter adds some power and fattens the bass.

Track 15.142: Bass 500 Hz Boost
A 7 dB boost at 500 Hz.

Track 15.143: Bass 1 kHz Boost
A 6 dB boost at 1 kHz adds some presence.

Track 15.144: Bass HSF 1 kHz Down
A –10 dB at 1 kHz on a high-shelving filter removes some attack and presence.

Plugin: Cubase StudioEQ

Acoustic guitar

Acoustic guitars can play various roles in the mix. On some mixes, often sparse ones, the guitar is one of the main instruments, usually along with the vocals. In such circumstances, we often want the guitar sound to be rich and full-bodied. On other mixes, the acoustic guitars only provide a reinforcement of the harmony and rhythm, which means that their body is less important. Very commonly acoustic guitars are treated with HPFs. The cutoff frequency is set depending on how much space there is, often starting from around 80 Hz—just below the lower E (subject to standard tuning). In many commercial mixes, the body of the guitar is fully removed, leaving it playing more of a rhythmical role than a harmonic one. When appropriate, the cut-off frequency is set so high that all that is left is a very thin, somewhat metallic sound—the played chords can hardly be discerned.

Another aspect of equalizing an acoustic guitar involves treating its body resonance, which could result in the unwanted accent of a specific spectral area or an increased level and sustain of specific notes. Finger squeaks are another thing we might want to reduce (or accent) using equalization. We can also control the timbre of the guitar by accenting or easing second- and third-order harmonics, which are found on the low-mids. The very highs of an acoustic guitar can be very sensitive to boosts, and even a small push from 10 kHz and above can make the sound appear thin and cheap. For this reason, if an acoustic guitar needs to be brightened, it could benefit from a wide-Q bell, rather than a shelving EQ.
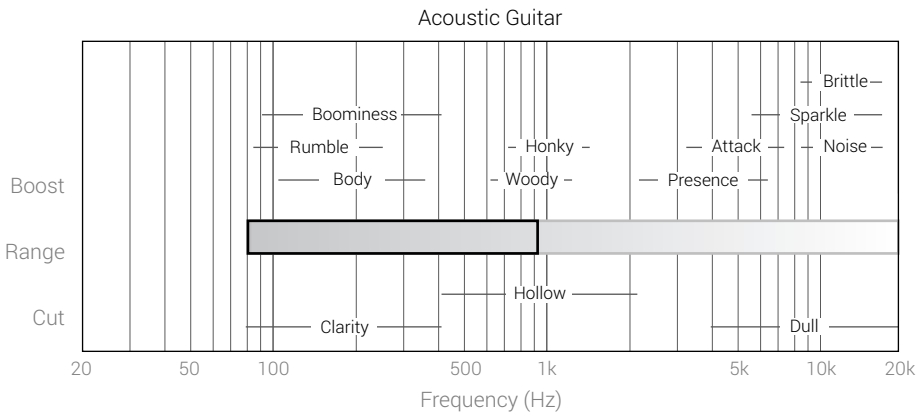
Figure 15.41 The frequency range of an acoustic guitar and relevant frequency ranges.

> **Track 15.145: aGtr HPF 220 Hz**
> The acoustic guitar in Track 15.85, with a 12 dB/oct HPF at 220 Hz. This is the source track to be equalized in the following samples.
>
> **Track 15.146: aGtr 300 Hz Boost**
> A 4 dB boost at 300 Hz adds some body, but also a hint of boominess.
>
> **Track 15.147: aGtr 1 kHz Boost**
> The honkiness on this track is the result of a 7 dB boost at 1 kHz.
>
> **Track 15.148: aGtr 4 kHz Boost**
> A 2 dB boost with a wide Q at 4 kHz is sufficient to add noticeable presence.
>
> **Track 15.149: aGtr HSF 7 kHz Up**
> This is the result of a 3 dB boost at 7 kHz on a high-shelving filter.
>
> Plugin: PSP MasterQ

Clean electric guitar

Electric guitars are an inseparable part of rock music, and on many productions provide the harmonic backbone. Most of them arrive at the mixing stage already equalized—whether by the tone controls on the guitar itself or those on the amplifier. Often our aim with clean electric guitars is making them well defined, yet without masking other instruments. The way the guitar is played is a factor in the way it is mixed. Some involve strumming chords while others are played in a phrasal fashion; some are played more rhythmically while others more melodically. The latter require more attention for spectral balance between different notes. Often we get more than one guitar, and, unless different on the recordings, we can distinguish them using equalization. Although to a lesser degree than with acoustic guitars, electric guitars can also benefit from low-end filtering, which is dependent on the density of the arrangement.
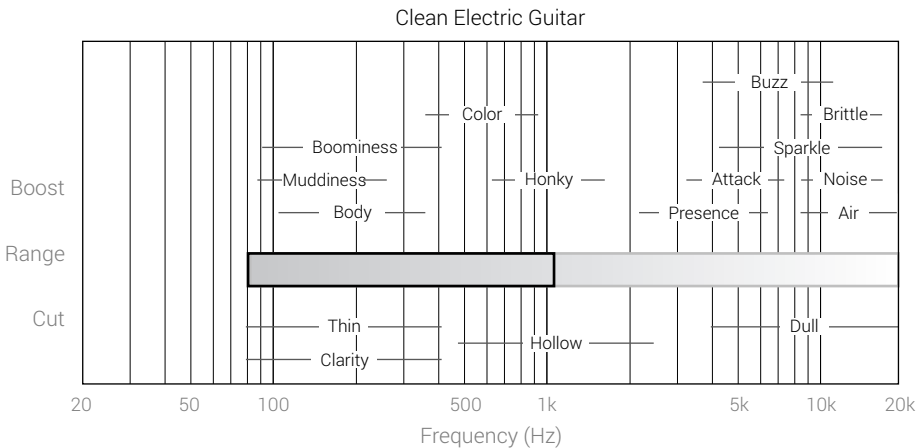
Figure 15.42 The frequency range of a clean electric guitar and relevant frequency ranges.

Track 15.150: cGtr No EQ
This source track is equalized in the following samples.

Track 15.151: cGtr HPF 230 Hz
A 6 dB/oct HPF at 230 Hz removes some of the guitar's body, but in mix context this type of treatment could yield clarity.

Track 15.152: cGtr 350 Hz Dip
This track is the result of a 3 dB attenuation at 300 Hz.

Track 15.153: cGtr 1 kHz Boost
Again, the honkiness on this track is the result of a 7 dB boost at 1 kHz.

Track 15.154: cGtr 4 kHz Boost
A 4 dB boost at 4 kHz adds some presence and attack.

Track 15.155: cGtr HSF 5 kHz Up
A 5 dB boost at 5 kHz on a high-shelving filter adds some spark.

Plugin: McDSP FilterBank E6

Distorted guitar

Perhaps the frequency-richest instrument is the distorted guitar. This is both a disadvantage and an advantage. It is a disadvantage since distorted guitars are masking animals—most of them have dominant energy from the very lows to the very highs, and they can easily cloud everything else. It is an advantage since we can shape their sound in many fascinating ways. Distorted guitars are super-EQ-sensitive. A boost of 3 dB on the high-mids can easily make an EQ whistle. They are responsive all the same to small level cuts. A quick listen to commercial tracks reveals that many distorted guitars are elegantly powerful, but not pushily powerful. Put another way, they are not just thrown into the mix at a loud level, but crafted more wisely to blend with other instruments. In many cases, distorted guitars also involve some stereo effects.
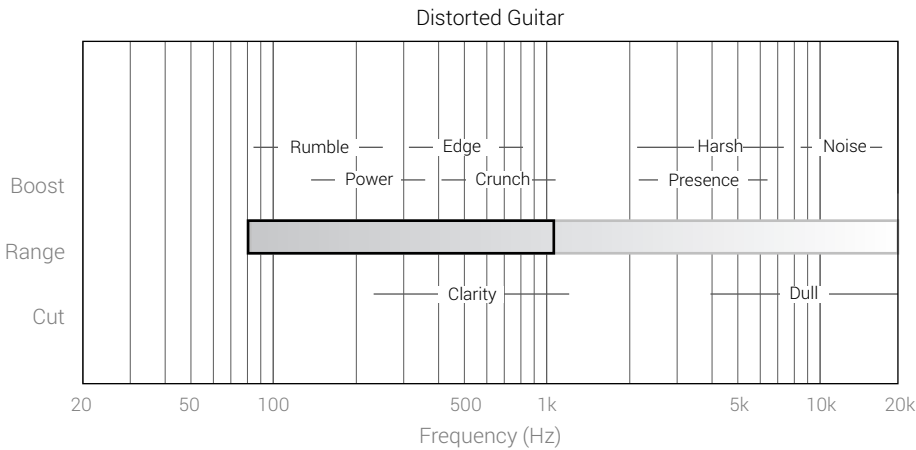
Figure 15.43 The frequency range of a distorted guitar and relevant frequency ranges.

Low-end rumble is often removed or attenuated. How much and up to what frequency the filter goes is done by ear. It is worth remembering that we can first clean the mix (going a bit higher with the filter) and later add a few more lows if we still feel these are missing. Much of the sound of distorted guitars has to do with the actual distortion being used, and there are as many flavors of guitar distortions as there are notes on a guitar. Often the high end contains grainy noise that the mix can live without. The mids can also benefit from softening, which generally clears some space for all the other instruments. But taking too much of the mids creates a very distant sound that is mostly associated with some metal subgenres. Another consequence of softening the mids is that the guitars become louder as the mix is played louder. This is the nature of any instrument with pulled mids, but it makes more sense with distortion guitars since they are more controlled in quiet levels but emphasized at louder levels (where things tend to get messier anyway).

Track 15.156: dGtr Source
This track is a combination of an HPF (115 Hz, 24 dB/oct), an LPF (10.7 kHz, 12 dB/oct), and a high-shelving filter (11.5 kHz, −11 dB), all applied to Track 15.81. This is the source track to be equalized in the following samples.

Track 15.157: dGtr 200 Hz Dip
A 6 dB attenuation at 200 Hz. This attenuates even further some low-frequency rumble.

Track 15.158: dGtr 500 Hz Dip
This track is the result of a 6 dB attenuation at 500 Hz.

Track 15.159: dGtr 1500 Hz Boost
A 4 dB boost at 1.5 kHz.

Track 15.160: dGtr 4 kHz Dip
A 4 dB attenuation at 4 kHz.

Plugin: MOTU MasterWorks EQ

16 Introduction to dynamic range processors

Chapters 17–21 deal with compressors, limiters, gates, expanders, and duckers. All fall into a group called *dynamic range processors*. The similarities between them are worth discussing as a background to the succeeding chapters.

Dynamic range

Dynamic range is defined as the difference between the softest and loudest sounds a system can accommodate. Our auditory system, digital systems (such as an audio sequencer), and analog systems (such as a loudspeaker) all have their dynamic range. Absolute level measurements are expressed using various dB scales—dBSPL, dBu, dBFS, etc. Dynamic range, like any *difference* between two level measurements, is expressed in plain dB. The table below shows the dynamic range of various systems.

System	*Approximate dynamic range (dB)*
Software mixer (32-bit float)	144
Best A/D converters available	122
Human hearing	120
Dynaudio Air 25 Active Monitors	113
Neumann *U87* Condenser Microphone	110
CD, or 16-bit integer audio file	96
Tape with noise reduction	75
FM radio	65
Vinyl	60

Say the quietest moment of an orchestral piece involves a gentle finger tap on the timpani, which generates 10 dBSPL. Then at the loudest moment the orchestra emits 100 dBSPL. That is a musical piece with 90 dB of dynamic range. A *U87* microphone placed in the venue would theoretically be able to capture the full dynamic range of the performance, including both the quietest and the loudest moments (although during the quietest moments the microphone noise might become apparent, and during the climax the microphone might distort). All but the cheapest A/D converters can accommodate 90 dB of dynamic range, so when the microphone signal is converted to digital and recorded onto a computer the dynamic range of the performance is retained.

The full dynamic range will also remain if the performance is saved onto a 16-bit integer file, or played through good studio monitors.

However, we will encounter a problem cutting this performance on vinyl—we have to fit 90 dB of dynamic range into media that can only accommodate 60 dB. There are a few things we can do. We can just trim the top 30 dB, so every time the orchestra becomes loud it is muted—an idea so senseless that no processor is designed to operate this way. We can just trim the bottom 30 dB of the performance, so every time the orchestra gets really quiet there will be no sound. Also not the wisest idea, although we do something similar every time we bounce from our audio sequencer (32-bit float) into a 16-bit file—we trim the bottom 9 bits (reducing from 144 to 96 dB).

How much is 60 dB? Readers are welcome to experiment: use pink noise from a signal generator, and bring it down by 60 dB using a fader. Chances are that at –60 dB you would not hear it. Now boost your monitor level slightly in order to hear the quiet noise, and then slowly bring the fader back up to 0 dB. Chances are you will find the noise fairly loud, and you might not even fancy bringing the fader all the way up; 60 dB is regarded in acoustics as the difference between the loud and the inaudible. We will discuss this further in Chapter 24 when we look at reverbs. The 60 dB experiment might be disappointing for some—you might wonder where half of our 120 dB hearing ability went. The auditory 120 dB was devised from what was thought to be the quietest level we can hear—0 dBSPL—and the threshold of pain, which is around 120 dBSPL. In practice, most domestic rooms have ambiance noise at roughly 30 dBSPL; 120 dBSPL is the level you'll be exposed to when standing a meter from the speaker in an extremely loud rock concert or a dance club. Under such circumstances, you might not be able to hear someone shouting next to you. A shouting voice is roughly 80 dBSPL. We perceive 90 dBSPL as twice as loud, which is already considered very loud. If there is such a thing as the dynamic range of everyday life, it is somewhere between 30 and 70 dBSPL—40 dB of dynamic range. Still, a studio control room may feature a dynamic range from 20 to 110 dBSPL.

Having said that, we might still want our vinyl to retain both the loudest and softest moments of our orchestral performance, and trimming would not let us do so. What we can do is **compress** the dynamic range. We can make loud levels quieter—known as **downward compression**; or we can make soft levels louder—**upward compression**. By way of analogy, it is like squeezing a foam earplug before fitting it into our ear—all the material is still there, but compressed into smaller dimensions.

Say we have the opposite problem: we have an old vinyl that we want to transfer onto a 16-bit audio file. We might want to convert 60 dB of dynamic range into 96 dB. A process called **expansion** enables this. We can make quiet sounds even quieter—**downward expansion**; or we can make loud sounds even louder—**upward expansion**. This is similar to what happens when we take the earplug out—it returns to its original size.

By and large, downward compression is the most frequent type of compression in mixing, and so the word downward *is omitted by convention—a compressor denotes a downward compressor. It is the same with expansion—downward expansion is more common, thus "downward" is omitted. This book follows these conventions.*

Now how does this relate to mixing? Well, very little … but at the same time, quite a lot. It is the mastering engineer's responsibility to fit the dynamic range of a mix onto different kinds of media (and they do compress vinyl cuts). Nowadays, with the decreasing popularity of tape machines, we hardly change between one system and another—during mixdown, we care little about the *overall* dynamic range of our mix. However, very often while we mix we want to make loud or quiet sounds louder or quieter. It is rarely the overall dynamic range we have in mind—what we really want to control is the **dynamics** of instruments.

Dynamics

The term *dynamics* in mixing is equivalent to the musical term—variations in level. Flat dynamics mean very little or next to no variations in level; vibrant dynamics mean noticeable variations; wild dynamics mean excessive fluctuations that are somewhat disturbing.

We distinguish between **macrodynamics** and **microdynamics**. In this book, **macrodynamics** are regarded as variations in level for events longer than a single note. We talk about macrodynamics in relation to the changing level between the verse and the chorus, or the level variations between snare hits, bass notes, or vocal phrases.

Microdynamics are related to level variations that happen within each note being played due to the nature of an instrument; for example, the attack and decay of a snare hit. We associate microdynamics with the **dynamic envelope** of sounds, which entail level variations that happen within each note (or hit) being played. The dynamic envelope constitutes the second half of an instrument's timbre (the first half being the spectral content). We can associate different envelopes with different instruments, as Figure 16.1 demonstrates. It is worth knowing that most instruments have some initial transient or level burst as part of their dynamic envelope. Very often, we employ dynamic range processors to control the microdynamics in these envelopes or reshape them to alter the instrument's timbre.
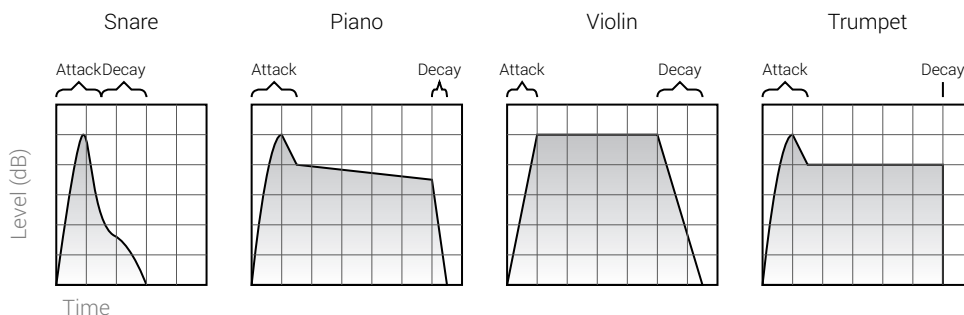
Figure 16.1 The typical dynamic envelope of various instruments. A snare has a quick attack caused by the stick hit; the majority of the impact happens during the attack stage, and there is some decay due to resonance. A piano also has an attack bump due to the hammer hitting the string; as long as the key is pressed, the sustained sound drops slowly in level; after the key is released, there is still a short decay period before the damping mechanism reaches full effect. The violin dynamic envelope in this illustration involves legato playing with gradual bowing force during the initial attack stage; the level is then kept at a consistent sustain level, and, as the bow is lifted, string resonance results in slow decay. In the case of a trumpet note, the initial attack is caused by the tongue movement; consistent airflow results in a consistent sustain stage, and as the airflow stops the sound drops abruptly.

📋 *It is worth noting the terminology used in Figure 16.1 and throughout this book. The attack is the initial level buildup and includes the fall to sustain level. Decay is used to describe the closing level descent. This is different from the terminology used for a synthesizer's ADSR (attack, decay, sustain, release) envelope.*

Dynamic range processors in a nutshell

Transfer characteristics

Say we regard an electrical wire as a processor, and its ends as the input and output. Practically speaking, signals entering the wire pass through unaffected. For any given input level, we get the same output level. We can say that the level ratio between the input and output is 1:1. Such a ratio is known as *unity gain*. We can draw a graph that shows the relationship between the input and output levels. Such a graph is called a *transfer characteristics graph*, or *input–output graph*, and is said to show the *transfer function* of a processor. It shows the input level on the horizontal axis and the output level on the vertical axis. The scale is given in decibels; 0 dB denotes the reference level of the system—the highest possible level of a digital system (which cannot be exceeded) or the standard operating level of an analog system (which can be exceeded). Figure 16.2 shows the transfer function of an electrical wire.
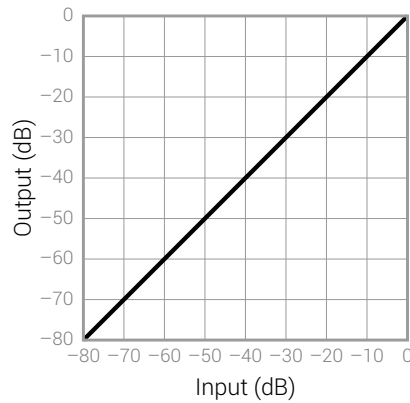
Figure 16.2 The transfer function of an electrical wire. There is no change in level between the input and the output—in other words, unity gain throughout.

As opposed to wires, dynamic range processors do alter the level of signals passing through them—the input and output levels might be different. We already know that a compressor reduces dynamic range, so we can expect input levels to be brought down. Figure 16.3 shows the transfer function of a compressor. We can see that the output's dynamic range is half that of the input.

If all dynamic range processors behaved like the compressor in Figure 16.3, they would all be rather limited (in fact, the earliest dynamic range processors were limited very much in this way). Treating all levels uniformly is rarely what we want, even in cases

where it is the whole mix we are treating. Most of the time, we only want to treat a specific level range—commonly either the loud or quiet signals. To enable us to treat selectively, all dynamic range processors allow us to limit treated levels and untreated levels using a parameter called **threshold**.

The threshold divides the full input level range into two sections. Depending on the processor, the treated level range might be above or below the threshold. Figure 16.4
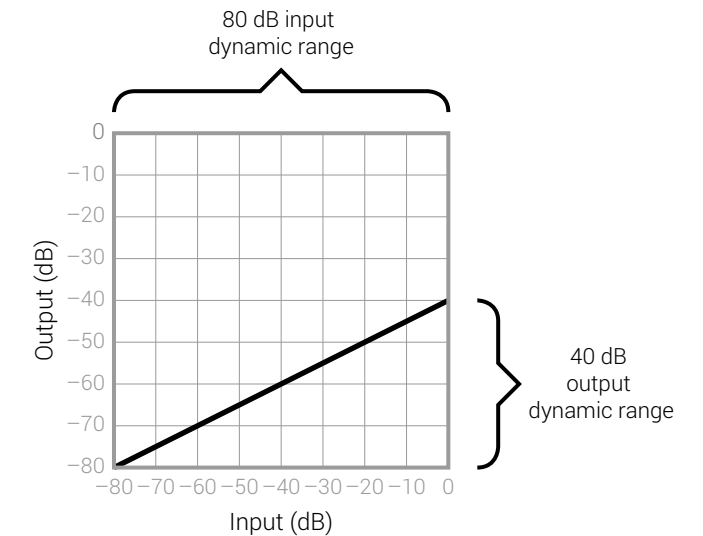
Figure 16.3 The transfer function of a compressor. The output dynamic range is half the input dynamic range.
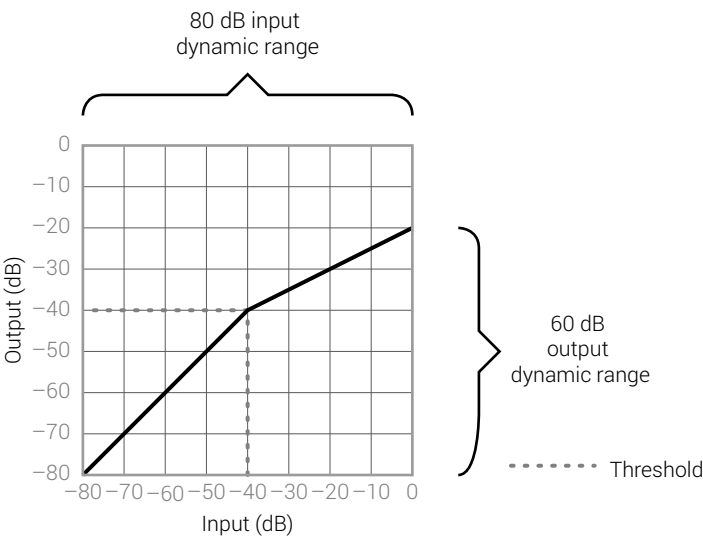
Figure 16.4 A compressor with threshold. The threshold is set to –40 dB. Levels below the threshold are not treated, and the input to output ratio is 1:1. Above the threshold, input signals are brought down in level. In this specific illustration, the input-to-output ratio above the threshold is 2:1.

shows a compressor with the threshold set at –40 dB. Being a downward compressor, loud levels are made softer, but only above the threshold—levels below the threshold are unaffected. Note that the output dynamic range has been reduced. It is also worth noting that the input-to-output ratio above the threshold is 2:1, resulting in the top 40 dB of input range (–40 to 0 dB) turning into 20 dB at the output (–40 to –20 dB). We will discuss ratios in greater detail in the following chapters.

The function of different processors

Figure 16.5 shows the transfer function of various dynamic range processors, and how levels are altered between the input and output. Here is a quick summary of what each of these processors does:

- **Compressor**—reduces the level of signals above the threshold, making loud sounds quieter.
- **Limiter**—ensures that no signal exceeds the threshold by reducing any signals above the threshold down to the threshold level.
- **Upward compressor**—boosts the level of signals below the threshold, making quiet sounds louder.
- **Expander**—reduces the level of signals below the threshold, making quiet sounds quieter.
- **Upward expander**—boosts the level of signals above the threshold, making loud sounds even louder.
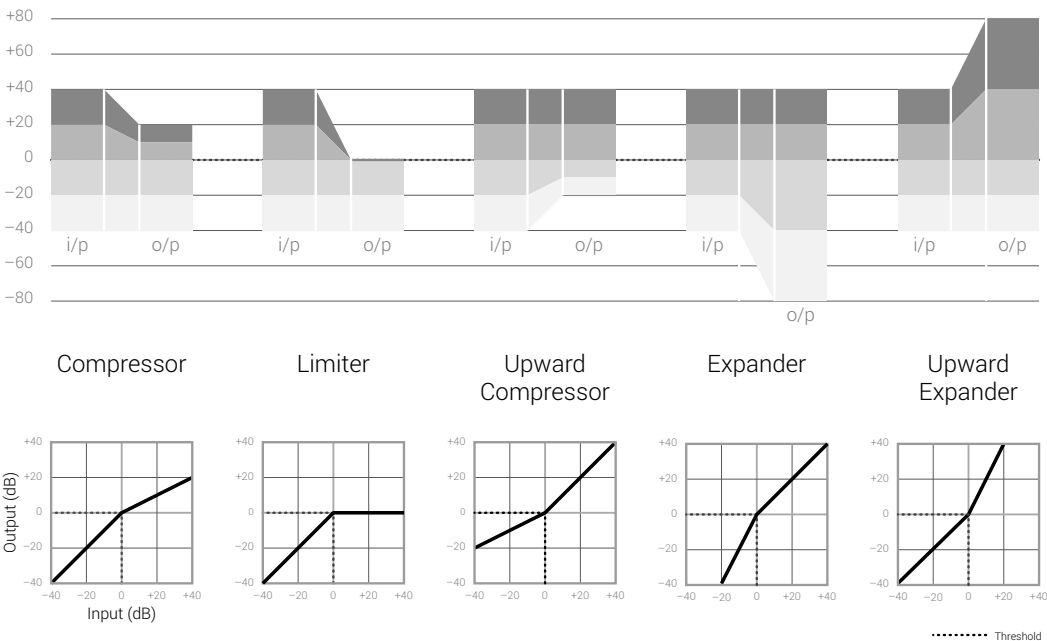
Figure 16.5 Sample transfer functions of compressors, limiters, and expanders, and their effect on input signals. The bottom graphs show the transfer function of each processor, all with a threshold set to 0 dB. The top row illustrates the relationship between the input and output levels.

All the processors above apply some ratio while transforming levels, making the level change dependent on the input level. For example, the compressor above has a 2:1 ratio above the threshold; +40 dB at the input is reduced to +20 dB (20 dB level change), while +20 dB is reduced to +10 dB (10 dB level change). The ratio of reduction is identical, but the level change is different.

Gates work on a slightly different principle—all signals below the threshold are reduced by a fixed amount known as range. With large range settings (say, –80 dB), signals below the threshold become inaudible, A ducker also reduces the signal level by a set range, but only does so for signals above the threshold. Figure 16.6 shows the function of a gate and a ducker, which can be summarized as follows:

- **Gate**—attenuates all signals below the threshold by a fixed amount known as range. Drastic attenuation makes everything below the threshold inaudible.
- **Ducker**—attenuates all signals above the threshold by a fixed amount—also known as range.
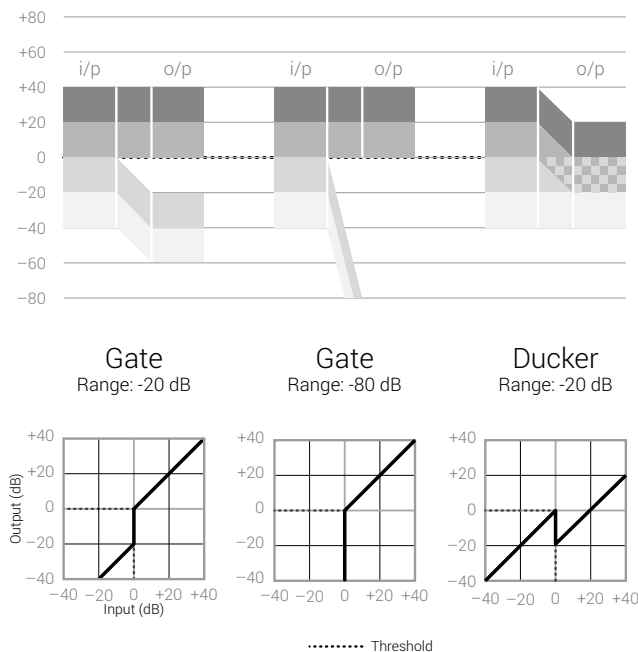
Figure 16.6 The transfer function of a gate with different range settings and a ducker. With small range settings, the gate simply attenuates everything below the threshold. With large range settings, the gate is said to mute signals below the threshold. A ducker attenuates signals above the threshold by a set range.

Pumping and breathing

Dynamic range processors can produce two artifacts known as pumping and breathing. **Pumping** is caused by quick, noticeable variations of level. We usually

associate pumping with loud level variations, such as those that can be the outcome of heavy compression or limiting. **Breathing** is the audible effect caused by varying noise (or hiss) levels. Often it is the quiet level variations that produce such artifacts, mostly due to the operation of gates or expanders.

Tracks 16.2 and 16.3 demonstrate pumping. The threshold on the compressor is set to –45 dB, the attack is set to slow, and the ratio is set to 60:1. In all of these tracks, the kick triggers the heaviest compression. It is worth noting the click on the very first kick, which is caused by the initial, drastic gain reduction.

Track 16.1: Original (No Pumping)

Track 16.2: Pumping (Fast Release)
The quick gain reduction and recovery is evident here; for example, the level of the hats fluctuates severely.

Track 16.3: Pumping (Medium Release)
The medium release causes slower gain recovery, which means levels fluctuate less. Although not as wild as in the previous track, the level changes are still highly noticeable.

Track 16.4: Pumping (Slow Release)
The slow release means that the gain hardly recovers before the next kick hits. The resultant effect cannot really be considered as pumping, since the gain fluctuations are too slow. However, it is hard to overlook how such slow recovery adds some dynamic movement to the loop.

Track 16.5: Breathing
The fluctuating noise level is associated with the term *breathing*. The vocal and the underlaid noise pass through a gate/compressor—both respond to the vocal. The gate causes the noise to come and go, while the varying amount of gain reduction on the compressor results in fluctuating noise level.

Plugin: Universal Audio GateComp

17 Compressors

Perhaps the most misused and overused tool in mixing is the compressor—which is especially worrying considering how predominant compressors are in contemporary mixes. It is no secret that compressors can make sounds louder, bigger, punchier, richer, or more powerful. Neither is it a secret that, since compressors were introduced per channel on analog consoles, the amount of compression in mixes has grown progressively and significantly—from transparent compression to evident compression, to heavy compression, to hypercompression, to ultracompression, so that by the start of the new millennium compression had been applied in so many mixes to such massive degrees that some people find it absurd:

• Are these VU meters broken?
• No.
• So why aren't they moving?
• Ah! It's probably my bus compression. Sounds like the radio, doesn't it?

If used incorrectly or superfluously, compressors suppress dynamics. Dynamics are a crucial aspect of a musical piece and a key messenger of musical expression. How impossible would our lives be if we could not alter the loudness of our voices? How boring would our singing be? A drummer would sound like a toy drum machine, and every brass instrument would resemble a car horn. If anything is to blame for the lifeless dynamics in contemporary music, it is the ultracompression trend.

A significant milestone in the story of compression was the release of Metallica's *Death Magnetic* in 2008. An album whose sound quality was so badly compromised by ultracompression—the hallmark of the infamous loudness war—that it didn't take a pair of trained ears to notice: the public complained. Subsequently, the term *loudness war* erupted from its oblivious boundary within the audio and music industries into the general public with help from national newspapers and social media. All because of one album; a decade or so too late in the view of many audio professionals.

Luckily, in recent years, more and more mixing and mastering engineers have left the ultracompression club and reverted back to more sensible compression. That's not to suggest that compressors should not be used, or used subtly. Many mixes still benefit from a hard-working compressor on many tracks. The real secret is to retain musical dynamics. We also have to distinguish the tightly controlled dynamics of pop music and the loss of dynamics caused by ultracompression. Some mixes present tightly

controlled dynamics, yet instruments sound alive. Ultracompression is not a requisite for commercial success; if anything, it is a threat.

One specific problem novice engineers have is that they do not know how the dynamics of a mix should sound. All they have for comparison is the dynamics of a mastered mix or, worse still, the dynamics of a mastered mix played through the radio. The latter can often sound squashed to death compared with a premastered mix. By trying to imitate the sound of a mastered album or the radio, the novice can easily limit the quality of the final master. Reviving a lifeless mix during mastering is hard, but if a mix has vibrant dynamics, mastering engineers can treat it as much as they want or are asked to.

One must pay attention to musical dynamics when compressing.

People often underestimate how hard compressors are to master. Making a beat punchy is not as easy as simply compressing it. There are numerous controls we can tweak and most of them are correlated. In order to obtain excellent results from a compressor, we must first understand how it works, the function of each control, and how these affect the dynamic behavior of each treated instrument. A compressor is such a powerful and versatile tool that some of its many applications are polar in nature (e.g., softening transients or emphasizing transients). Perhaps, after reading this chapter, those who used compressors in the past will feel that they have been flying a spaceship when they thought they were driving a car. It is a long chapter, but the knowledge provided can make the difference between the flourish techniques of the seasoned pro and the random trials of the novice. The consequences for the mix can be profound.

The course of history

One problem in early radio transmissions was that, if program levels exceeded a certain threshold, the transmitter would overload and blow—never a good thing, particularly during the climax of a thrilling sports broadcast. Radio engineers had to observe levels meticulously and ride the gain of live broadcasts so that levels did not exceed the permitted threshold (Figure 17.1). While these engineers did their best, only a fortune teller could predict sudden level changes, and it was safe to assume that no fortune teller would apply for the job. Even if they did, there is a limit to just how quickly one can respond to sounds or physically move a fader—sudden level changes would still be an issue. While **peak control** was aimed at protecting the transmitter, gain-riding was also performed in order to **balance** the level of the program, levelling the different songs and the presenter's voice.

By the 1920s, people such as James F. Lawrence Jr. (who later designed the famous LA-2A) already had ideas about building a device that would automate the gain-riding process. The concept was to feed a copy of the input signal to a side-chain, which, based on the incoming level, would determine the required amount of gain reduction. The side-chain was connected to a gain stage, which applied the actual gain reduction on the signal (Figure 17.2). The name given to these early devices was *leveling amplifier* (suggesting more of a balancing function) or *limiter* (suggesting more of a peak control function). However, early models had very slow response to sudden level changes, and thus they did not really limit the level of the signal—they behaved much
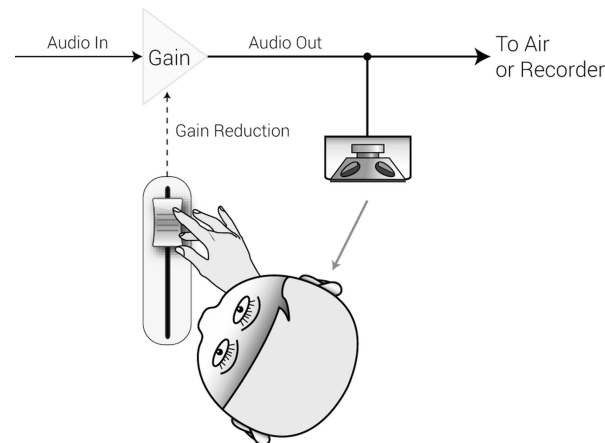
Figure 17.1 Manual gain-riding. The engineer gain-rides levels with relation to the speaker output. In this illustration, a VCA fader is employed.

(a)   (b)
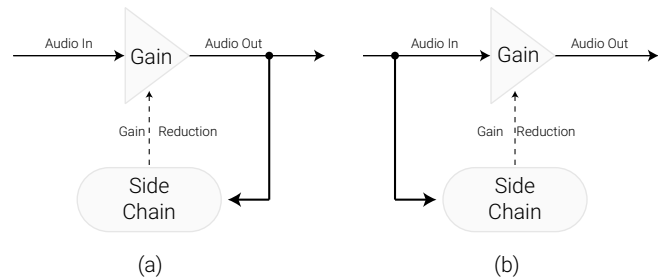
Figure 17.2 (a) A feedback-type compressor. The input into the side-chain is taken past the gain stage, which resembles the way manual gain-riding works. Early compressors benefited from this design as the side-chain could rectify possible inaccuracies of the gain stage (for example, when it did not apply the required amount of gain reduction). However, this design had more than a few limitations; for instance, it did not allow a look-ahead function. (b) A feed-forward type compressor has the input into the side-chain taken before the gain stage. As most modern compressors are based on this design, it will be discussed in the rest of this chapter.

like today's compressors. With advances in technology, true limiting was eventually possible, and the distinction between compressors and limiters had been made.

Studio engineers quickly borrowed compressors for their recording studios as, just like radio engineers, they had to gain-ride live performances. At the time, it was necessary to contain peaks so that when music was cut on the fly to disc it would not distort. Later, it was utilized to protect tapes from saturating, and later still to protect digital systems from clipping. Compressors were also employed to even out the dynamics of a performance; for example, when vocals change from soft to loud. By the 1960s, units such as the Urei 1176 LN, *Fairchild 670*, and LA-2A were already a common sight in control rooms.

Compressors alter both the dynamic envelope of the source material and, like any other nonlinear device, add some distortion—together we get the distinctive effect of

compression. The original intention of compressors was to alter the dynamic range while leaving as little audible effect as possible. However, it soon became apparent that the effect of compression, like that caused by an overloaded tape, can be quite appealing. One sonic pioneer who understood this was Joe Meek, who, instead of concealing the effect of compressors, used it as part of his distinctive sound. Among his many highly acclaimed recordings was the first British single ever to top the U.S. Billboard chart—"Telstar" by The Tornados (1962). Nowadays, the effect of compression can be heard in nearly every mix.

Sometimes we want the compressor to be transparent.
Sometimes we use it for its distinctive effect.

When it comes to the use of a compressor as a level balancer, the end to this brief history is somewhat interesting. The compression effect becomes more evident with heavier compression. If a vocal performance changes from crooning to shouting, the effect will be more noticeable on the shouting passage, being subject to more compression. While after compression the performance might appear to be consistent in level, the shouting passage could have a distinct compression effect that will be missing from the crooning passage. To combat situations such as this, manual gain-riding still takes place before compression (despite the latter being invented to replace the former).

The sound of compressors

No two compressor models sound alike, certainly no two analog compressors. Some can work better for drums and some can work better on vocals; some add warmth, some an extra punch; some can be transparent while others produce a very obvious effect. Each compressor has a character, and in order to have a character something must be different from its counterparts.

We can generalize what we want from a *precise* compressor; here are just a few points:

• We want compression to start at a consistent point with relation to the threshold.
• We want the gain stage to act uniformly on all signal levels and with consistent response times.
• We want accurate performance from the ratio function.
• We want to have the ability to dial attack and release times of our choice, even if these are very short.
• We want the attack and release envelopes to be consistent and accurate.

In practice, there is little challenge building a precise digital compressor. A quick online search will quickly yield a free algorithm for exactly such a compressor, but it won't have character. Analog designs have a character due to their lack of precision, and each compressor is inaccurate in its own unique way. How the ratio behaves or the nature of the attack and release functions defines much of the compressor sound. It is mostly these aspects that earned vintage models much of their glory. In order to introduce some character into digital compressors, designers have to choose where and how to introduce deviations from the precise design.

Figure 17.3 From top to bottom: The Chandler Limited RS124; The Empirical Lab's Distressor; The Focusrite Blue 330. The three have distinctively different sounds. Courtesy of Abbey Road Institute, London.

Figure 17.4 A compressor plugin: the Digidesign *DigiRack Compressor/Limiter Dyn 3*.

Principle of operation and core controls

The explanation of how compressors work is often oversimplified and inaccurate. To really understand how this popular tool works, it is imperative to look at its internal building blocks—consisting of a few simple stages. Figure 17.5 focuses on the main stages and the controls linked to each. The illustration is based on the
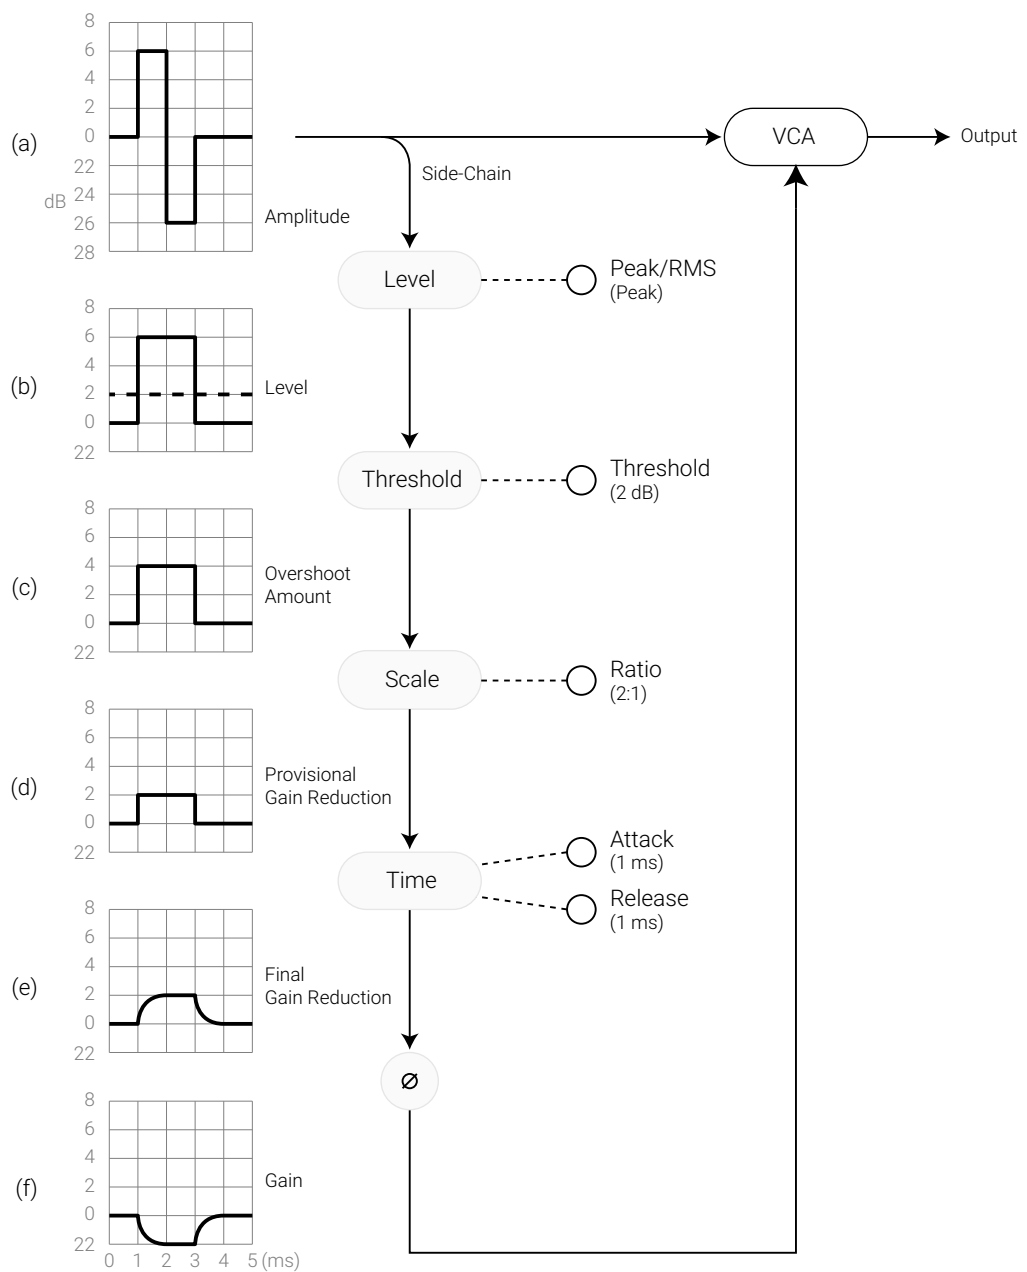
Figure 17.5 Inside a compressor. The vertical chain shows the main stages within the side-chain, and the controls link to each stage. The nature of the signal between one stage and another is shown to the left of the chain, along with a sample graph that is based on the settings shown under each control name.

modern VCA analog design—pretty much the blueprint for most modern compressors, including digital ones.

Gain

The gain stage is responsible for attenuating (and in some cases also boosting) the input signal by a set number of dBs (which is determined by the side-chain). The gain stage in Figure 17.5 is based on a VCA, but other types of gain stages may be used. It pays to learn the differences between them, if only because many of today's digital compressors try to imitate the sound of old analog ones:

- **Vari-mu**—the earliest compressor designs were based on variable-mu tubes (valves); essentially, variable gain amplifiers. Vari-mu designs have no ratio control; they apply incremental amount of gain reduction with relation to input levels (a behavior similar to soft-knee, which is explained later). But this only happens up to a point, when the compressor returns to linearity. This characteristic works well for percussive instruments as loud transients are not clamped down. Vari-mu designs have faster attack and release times than optical designs, but they are not as fast as VCA designs or FET.
- **FET**—as small transistors started replacing the large tubes, later compressor designs were based on field effect transistors (FETs). They offered considerably faster attack and release times than vari-mu and incorporated a new feature—ratio. Just like vari-mu designs, the compression ratio tends to return to linearity with a very loud input signal.
- **Opto**—the side-chain of an optical compressor controls the brightness of a bulb or LED. On the gain stage, there is a photo-resistive material, which affects the amount of applied gain. Much like our pupils takes a bit of time to contract and dilate in response to changing light intensity, photo-resistive material has slow response compared with musical dynamics; thus, optical compressors exhibit the slowest response times of all compressors. Moreover, their attack and release curves are less than precise (especially with older designs), giving these compressors a very unique character. Optical designs are known to produce a very noticeable effect, which many consider appealing.
- **VCA**—of all their analog equivalents, solid-state voltage-controlled amplifiers (VCAs) provide the most precise and controllable gain manipulation. Their native accuracy broadened the possibilities compressor designers had, and made VCAs the favorite in most modern designs.
- **Digital**—digital compressors work using a set of mathematical operations. When it comes to precision, a digital compressor can be as precise as it gets. Their response time can be instant, which means no constraints on attack and release times. Digital compressors can also offer perfectly precise ratio, attack, and release curves.

Level detection and peak vs. RMS

As the signal enters the side-chain, it first encounters the level stage. This is where its bipolar amplitude (Figure 17.5a) is converted into a unipolar representation of level (Figure 17.5b). At this point, the level of the signal is determined by its peak value. Instead of peak-sensing, it might be in our interest to have RMS-sensing, so the compressor responds to the loudness of incoming signals rather than to their peaks (vocals are often better compressed that way). An RMS function (or other averaging

function) might take effect at this stage. By way of analogy, it is as if we have replaced a peak meter within the level stage with a VU meter.

A compressor might support peak-sensing only or RMS-sensing only, or offer a switch to toggle between the two. Some compressors also allow us to dial a setting between the two. If no selection is given, the compressor is likely to be an RMS one.

Threshold

The threshold is the level above which gain reduction starts. Any signal exceeding the threshold is known as an *overshooting signal* and would normally be reduced in level. The more a signal overshoots, the more gain reduction is applied. Signals below the threshold are generally unaffected (we will look at two exceptions to this later). The threshold is most often calibrated in dB.

A threshold function on a compressor comes in two forms: a **variable threshold** or a **fixed threshold**. A compressor with variable threshold provides a dedicated control with which the threshold level is set. A compressor with fixed threshold (Figure 17.7) has an input gain control instead—the more we boost the input signal, the more it will overshoot the fixed threshold. This is similar to the way compression is achieved when overloading tapes. To compensate for the input gain boost, an output gain control is also provided. One advantage of the fixed threshold design (with analog units, anyway) is that the noise introduced by the gain stage is often attenuated by the output gain.
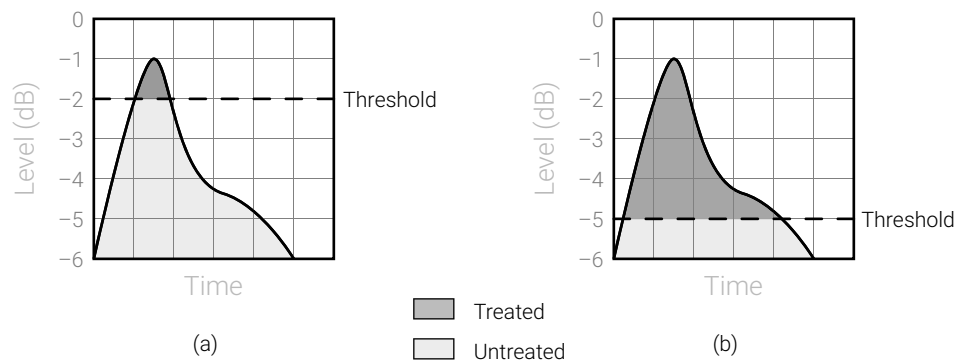
Figure 17.6 Threshold setting. (a) Higher threshold means that a smaller portion of the signal is treated. (b) Lower threshold results in a larger portion of the signal being treated.
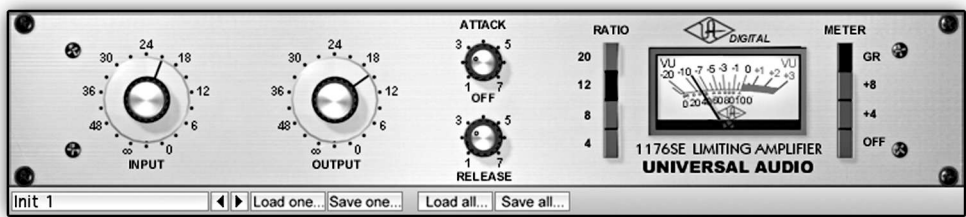
Figure 17.7 The Universal Audio *1176SE*. This plugin for the UAD platform emulates the legendary Urei 1176LN. It has fixed threshold and the degree of compression is partly determined by the input gain amount.

On variable threshold designs, the same noise is usually boosted by the output (or makeup) gain in order to compensate for the applied gain reduction.

Track 17.1: Dropping Threshold
This track involves a continuously dropping threshold, from 0 dB down to –60 dB at a rate of –7.5 dB per bar. The compressor is configured to bring down signals above the threshold as quickly as possible (fastest attack: 10 μs) and as much as possible (highest ratio: 100:1). Therefore, the lower the threshold, the larger the portion of the signal being brought down in level, resulting in a gradual decrease in overall level.

Track 17.2: Dropping Threshold Compensated
This track is similar to the previous one, except after compression has been applied the track level is gradually boosted to compensate for the loss of level caused by the compressor. This reveals a better increasing compression effect over time.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3
Drums: Toontrack EZdrummer

As can be seen in Figure 17.5c, the threshold stage is fed with the level of the side-chain signal. The output of this stage is the *overshoot amount*, which indicates by how much an overshooting signal is above the threshold. For example, if the threshold is set to 2 dB and the signal level is 6 dB, the overshoot amount is 4 dB. If the signal level is below the threshold, the overshoot amount is 0 dB.

Ratio

Ratio can be compared with gravitational force, which determines the extent to which objects are pulled toward the ground. There is less gravity on the moon, so astronauts can jump higher. Ratio determines the extent to which overshooting signals are reduced toward the threshold. The lower the ratio, the easier it is for signals to jump above the threshold. Figure 17.8 shows the effect of different ratios on an input signal.
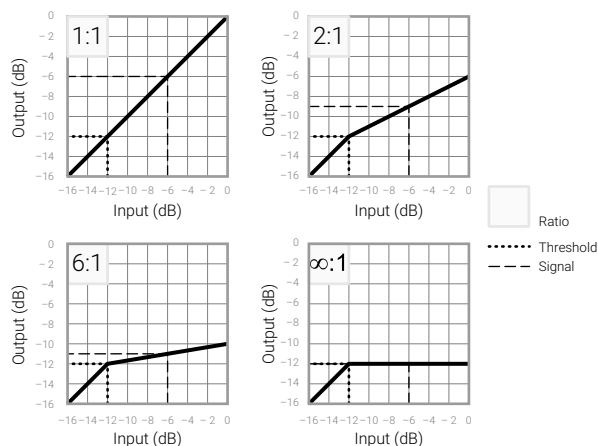
Figure 17.8 The effect of different ratios on the level envelope of a waveform.
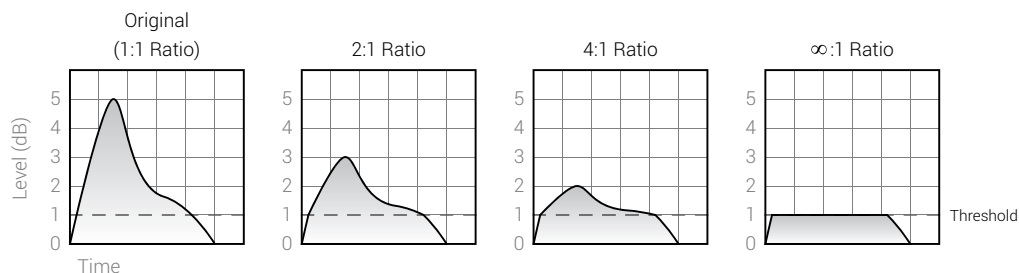
Figure 17.9 Different ratios on transfer characteristics graphs. In all these graphs, the threshold is set to –12 dB and the input signal is –6 dB (6 dB above the threshold). With a 1:1 ratio, the output signal leaves the compressor at the same level of –6 dB (6 dB above the threshold); with a 2:1 ratio, the signal leaves at –9 dB (3 dB above the threshold); with a 6:1 ratio, the signal leaves at –11 dB (1 dB above the threshold); and with ∞:l, the signal leaves at the same level as the threshold.

Newtonian physics aside, once the signal overshoots the threshold, the ratio control determines the ratio between input and output level changes (as the input-output notation suggests). For example, with a 2:1 ratio, an increase of 2 dB above the threshold for input signals will result in an increase of 1 dB above the threshold for output signals. The ratio determines how overshooting signals are scaled down. A 1:1 ratio (unity gain) signifies that no scaling takes place—a signal that overshoots by 6 dB will leave the compressor 6 dB above the threshold. A 2:1 ratio means that a signal overshooting by 6 dB is scaled down to half of its overshoot amount and leaves the compressor at 3 dB above the threshold. A 6:1 ratio means that a signal overshooting by 6 dB is scaled down to a sixth and leaves the compressor at 1 dB above the threshold. The highest possible ratio is ∞:1 (infinity to one), which results in any overshooting signal being trimmed to the threshold level. Often, a ratio of ∞:1 is achieved by using an extremely high ratio, such as 1,000:1. Figure 17.9 illustrates these four scenarios.

Track 17.3: Raising Ratio
This track involves a threshold at –40 dB and a continuously rising ratio from 1:1 up to 64:1. The ratio doubles per bar, starting with 1:1 at the beginning of the first bar, 2:1 at the beginning of the second bar, 4:1 at the beginning of the third bar, and so on. The higher the ratio, the harder overshooting signals are reduced toward the threshold, resulting in a gradual decrease in the overall level.

Track 17.4: Raising Ratio Compensated
This track is identical to the previous one, only that the loss of level caused by the compressor is compensated using a gradual gain boost. Here again, the increased compression effect over time is easily discerned.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3
Drums: Toontrack EZdrummer

One characteristic of vintage compressors, notably vari-mu and FET designs, is that the ratio curve maintains its intended shape, but morphs back to unity gain with loud
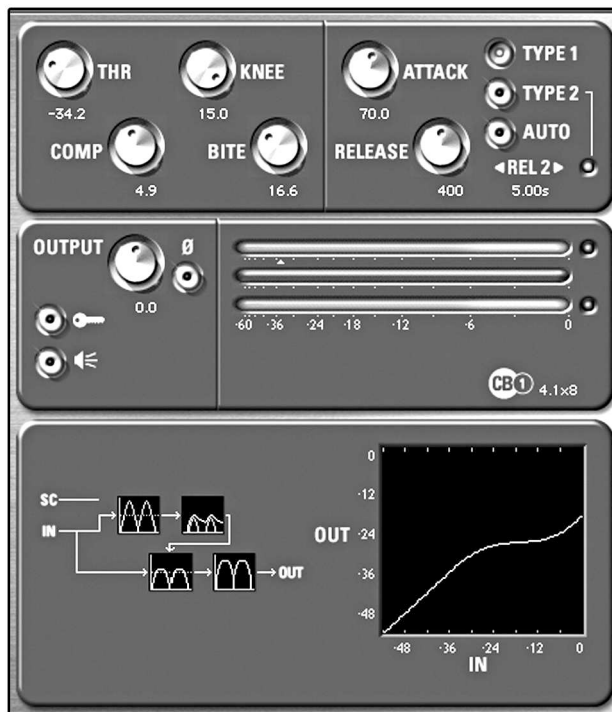
Figure 17.10 The *McDSP CB1* plugin. The input-output plot shows that compression starts around the threshold and then develops into heavy limiting. However, around –12 dB, the ratio curve changes course and starts reverting to unity gain. Ratio slopes of this nature, which deviate from the textbook-perfect shape, can provide a characteristic sound that is usually a trait of vintage analog designs.

signals. This type of ratio curve tends to complement transients, as the very loud peaks are not tamed—compression is applied, but less dynamics are lost. Drum heaven. Many digital compressors imitate such a characteristic behavior by employing similar ratio curves. The McDSP *CB1* in Figure 17.10 is one of them.

The ratio is achieved within a compressor by the scale stage (Figure 17.5). The scale stage is the operational heart of the compression process—this is where the overshoot amount is converted into the provisional gain reduction. The output of the scale stage is nearly fit to feed the gain stage. The ratio essentially determines the percentage by which the overshoot amount is reduced to 0 dB. A 1:1 ratio means that the overshoot amount is fully reduced to 0 dB, which means that no gain reduction takes place. With a 2:1 ratio, the overshoot amount is reduced by half. For example, 4 dB of overshoot becomes 2 dB.

Attack and release

Modern compressors can respond to sudden level changes instantly, but such a quick response is not always sought. For example, in order to retain some of the instrument's natural attack, we often want to let some of the initial level burst through the

compressor unaffected (or lightly affected). In order to do so, we need to be able to slow down the compressor response times. Similarly, if a healthy amount of gain reduction drops too fast, the gain recovers too quickly, which may result in pumping. To prevent this, we need a way of controlling the rate at which gain reduction drops.

The attack and release are also known as *time constants* or *response times*. The attack determines how quickly gain reduction rises, while release determines how quickly gain reduction falls. Essentially, a longer setting on either will simply slow down the rate at which gain reduction increases (attack) or decreases (release). For example, in Figure 17.5d, the gain reduction rises instantly from 0 to 2 dB and then falls back instantly to 0 dB. Having 1 ms of both attack and release means, in this case at least, that it takes 1 ms for the gain reduction to rise and fall (Figure 17.5e).

Both the attack and release times are typically set in milliseconds. Attack times usually range between 0.010 ms (10 μs) and 250 ms. Release times are often within the 5–3,000 ms (3 seconds) range. It is important to understand that both times determine how quickly the gain reduction can change, and not the time it takes it to change. In practice, both define how long it takes the gain reduction to change by a set number of dB. For example, 1 second of release time might denote that it takes the gain reduction 1 second to drop by 10 dB. Accordingly, it would take it half a second to drop by 5 dB. This behavior can be compared with the ballistics of a VU meter. A quick peak drop would appear gradual on a VU meter. When peaks drop slowly, the VU reading will be very similar to that of a peak meter. The attack and release have a very similar effect on gain reduction.

Figure 17.11 shows the effect of different attack and release settings on a waveform. On all graphs, the ratio is 2:1, and the threshold is set to 6 dB. The original input signal (a) rises instantly from 6 to 12 dB, and then drops instantly back to 6 dB. In all cases, the overshoot amount is 6 dB, so with the 2:1 ratio the full gain reduction amount is 3 dB. We can see in (b) that if there is no attack and no release, the overshooting signal is constantly reduced by 3 dB. When there is some attack and release (c–e), it takes some time before full gain reduction is reached and then some time before gain reduction ceases. It is worth noting what happens when the release is set and the original signal drops from 12 dB back to 6 dB. Initially, there is still 3 dB of gain reduction, so the original signal is still attenuated to 3 dB below its original level. Slowly, the gain reduction diminishes and only after the release period has passed does the original signal return to 6 dB. This is the first exception for signals below the threshold being reduced in level. The gain reduction graph below (d) helps us to understand why this happens.

Some compressors offer an **auto attack** or **auto release**. When either is engaged, the compressor determines the attack or release times automatically. Mostly, this is achieved by the compressor observing the difference between the peak and RMS levels of the side-chain signal. It is worth knowing that in auto mode, neither the attack nor the release is constant (like when we dial the settings manually). Instead, both change in relation to the momentary level of the input signal. For example, a snare hit might produce a faster release than a xylophone note. Thus, auto attack and release do not provide less control—they simply provide an alternative kind of control. Auto release, on respected compressors at least, has a good reputation.
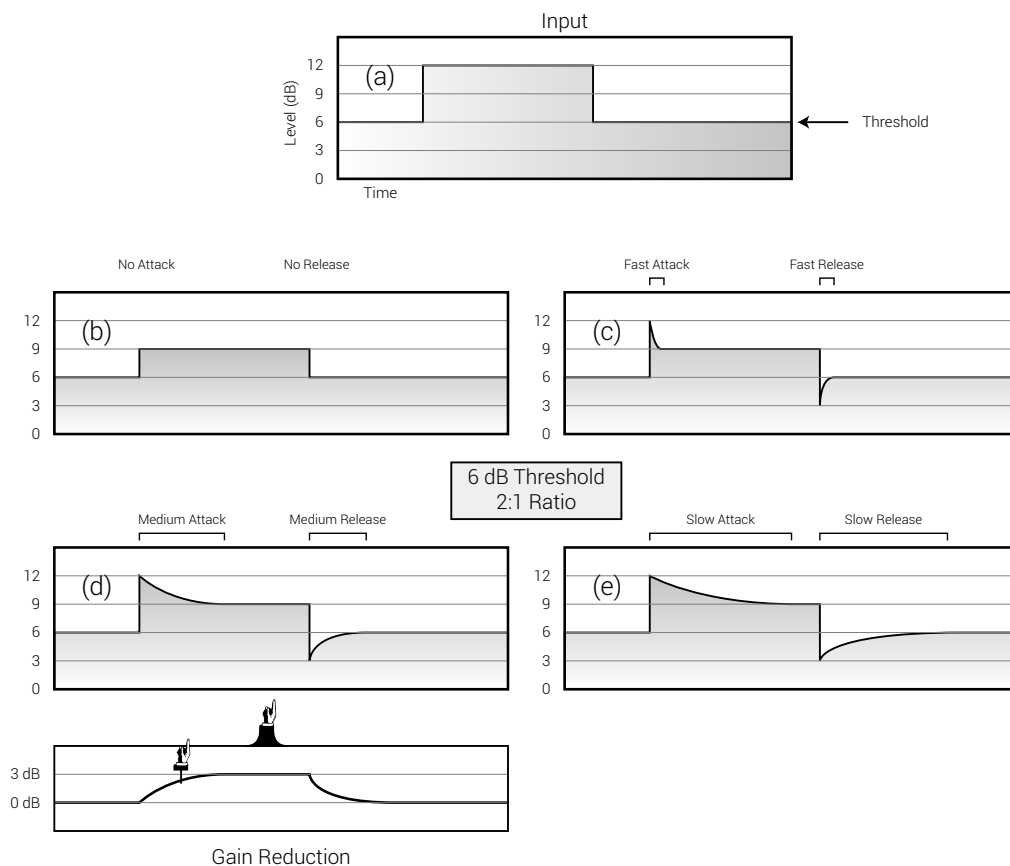
Figure 17.11 The effect of different attack and release settings on a waveform. (a) The original levels before compression. (b–e) The resultant levels after compression.

Track 17.5: Noise Burst Uncompressed
This track involves an uncompressed version of a noise burst. First, the noise rises from silence to –12 dB, then to –6 dB, and then it falls back to –12 dB and back to silence. The compressor in the following tracks was set with its threshold at –12 dB and a high ratio of 1,000:1. Essentially, when full gain reduction is reached, the –6 dB step should be brought down to –12 dB.

Track 17.6: Noise Burst Very Fast TC
With the attack set to 1 ms and the release to 10 ms, we can hardly discern changes in level caused by the attack and release functions. Still, there is quick chattering when the level rises to –6 dB and when it falls back to –12 dB.

Track 17.7: Noise Burst Fast TC
This track involves an attack time set to 10 ms and release to 100 ms. The click caused by the attack can be heard here, and so can the quick gain recovery caused by the release.

> **Track 17.8: Noise Burst Medium TC**
> An attack of 25 ms and 250 ms of release. A longer attack effect can be heard, and the gain recovery resulting from the release is slower.
>
> **Track 17.9: Noise Burst Slow TC**
> The most evident function of the time constants is achieved using slow attack and release: 50 and 500 ms, respectively. Both the drop in noise level resulting from the attack and the later rise resulting from the release can be clearly heard on this track.
>
> Plugin: Sonnox Oxford Dynamics

Much of a compressor's character (or any other dynamic range processor, for that matter) is determined by the attack and release functions; in particular, their **timing laws**. The timing laws determine the rate of change the attack and release apply on gain reduction, and in their most simple form these can be either exponential or linear. There is some similarity here to exponential and linear fades, except that fades are applied on the signal itself, while the attack and release are applied on gain reduction. Generally speaking, exponential timing laws tend to sound more natural. Linear timing laws tend to draw more color and effect, which is often associated with the sound of some favorite analog units. Very few compressors give us control over the compressor's timing laws; the compressor of the Sonnox *Oxford Dynamics* in Figure 17.13 is one of the few that does (the "NORMAL" field above the plot denotes exponential law).

As seen in Figure 17.5, both the attack and release are controls linked to the time stage. The input to the time stage is the provisional amount of gain reduction. The time stage slows down sudden changes to that gain reduction. As long as the input to the time stage is higher than its output, the gain reduction will keep rising at the rate set by the attack. The moment the input is lower than the output, the gain reduction starts to drop at the rate set by the release. Figure 17.12 illustrates this.
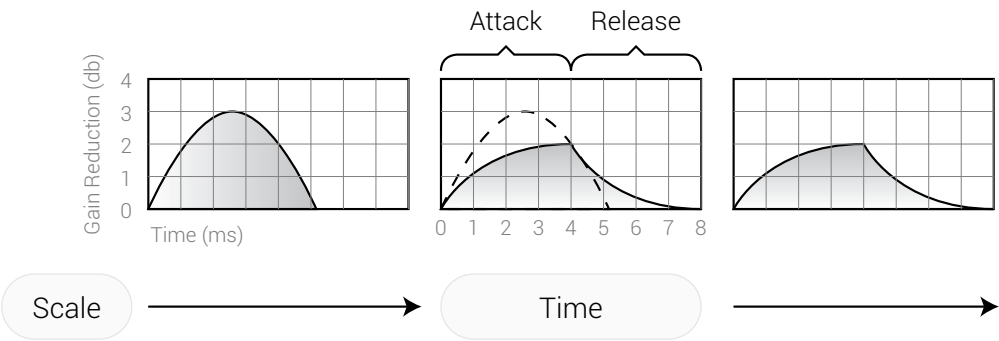
Figure 17.12 The input to the time stage arrives from the scale stage and denotes the provisional amount of gain reduction. The attack within the time stage slows down the gain reduction growth. Note that even when the input to the time function starts to fall (2.5 ms), the output still rises as the provisional amount of gain reduction is still higher than the applied amount. Only when the provisional amount falls below the applied amount (4 ms) does the latter start to drop. At this point, the attack phase ends and the release phase starts.

The following tracks demonstrate the differences between exponential and linear timing laws. It should be obvious that in linear mode, the compression effect is more evident. It produces noticeable distortion. Although maybe not to such an extreme degree, this characteristic sound is sometimes what we are after.

Track 17.10: Metal Uncompressed

Track 17.11: Metal Exponential Mode

Track 17.12: Metal Linear Mode

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack EZdrummer

One crucial thing to understand is that attack and release are applied on gain reduction, and that the time stage, where these are applied, is unaware of the threshold setting. Many sources state incorrectly that the release only occurs when the signal level drops below the threshold. In reality, both the attack and release affect gain reduction (and in turn the signal level), even when the signal level changes above the threshold. Figure 17.13 demonstrates this.
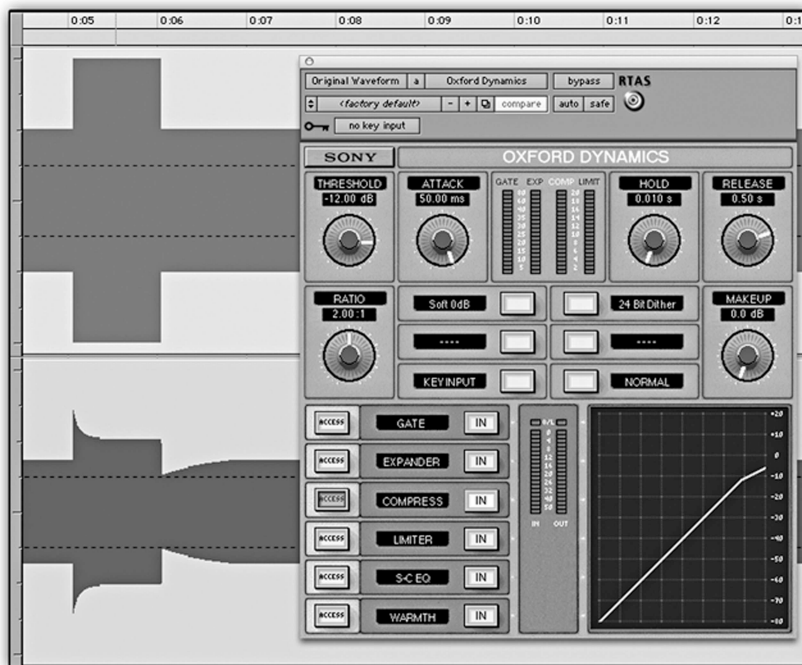
Figure 17.13 Attack and release above the threshold. The top waveform rises from –6 to 0 dB, and then falls back to –6 dB. This waveform passes through an *Oxford Dynamics* plugin with the shown settings, and the post-compression result is the bottom waveform. The compressor threshold is set to –12 dB (dashed lines) so all the level changes in the top waveform happen above the threshold. On the bottom waveform, which shows the output of the compressor, we can see the action of both the attack and release. Analog compressors behave in the same way.

Despite what many sources suggest, the release function is not related to
the signal dropping below the threshold.

The top waveform in Figure 17.13 is known as a noise burst, and is commonly used to demonstrate the time function of dynamic range processors. However, it could be just as easily demonstrated using a voice rising from levels already above the threshold and then falling back to levels above the threshold (something like ah-Ah-AH-Ah-ah, where only the ah is below the threshold). The fact that the attack and release affect level changes even when these happen above the threshold means that the dynamics are controlled to a far greater extent than they would be if compressors only acted when signals overshoot or drop below the threshold. The attack and release would still be applied on the AH just like on the Ah. Every so often, vocals do fluctuate in level after overshooting the threshold, and the release and attack still affect our ability to balance them. Also, attack and release will still act on signals fluctuating above the threshold, even when it is set very low.

Track 17.13: Noise Burst Low Threshold
A compressed version of Track 17.5, with the threshold set to –40 dB and ratio to maximum. You should hear the first attack when the noise rises to –12 dB, and then again when it rises to –6 dB (which happens above the threshold). When the noise falls back to –12 dB (a level above the threshold), the gain recovery caused by the release function is evident.

The following two tracks involve a threshold at –50 dB and a ratio of 1.5:1. Apart from the very first 300 ms and the closing silence, the vocal is always above the threshold. The only differences between the two tracks are the attack and release settings. There is no doubt that the two are distinctly different, demonstrating that both the attack and release affect level variations above the threshold.

Track 17.14: Vocal Compression I

Track 17.15: Vocal Compression II

Plugin: Sonnox Oxford Dynamics

Hold

Very few compressors provide a hold parameter, which is also linked to the time function. In simple terms, hold determines how long gain reduction is held before the release phase starts. Figure 17.14 illustrates this. In practical terms, the function of hold on a compressor is achieved by altering the release rate, so during its early stages gain reduction hardly changes. Although the results are only similar to those shown in Figure 17.14, the overall impression is still as if gain reduction is held.

Phase inverse

The phase stage in Figure 17.5 is shown for explanation's sake—it might not be part of a real-world compressor. The output of the time stage is the final gain reduction, and it involves a positive magnitude (Figure 17.5e). What the gain stage expects is the amount of *gain* that should be applied on the signal, not the amount of *gain reduction*.
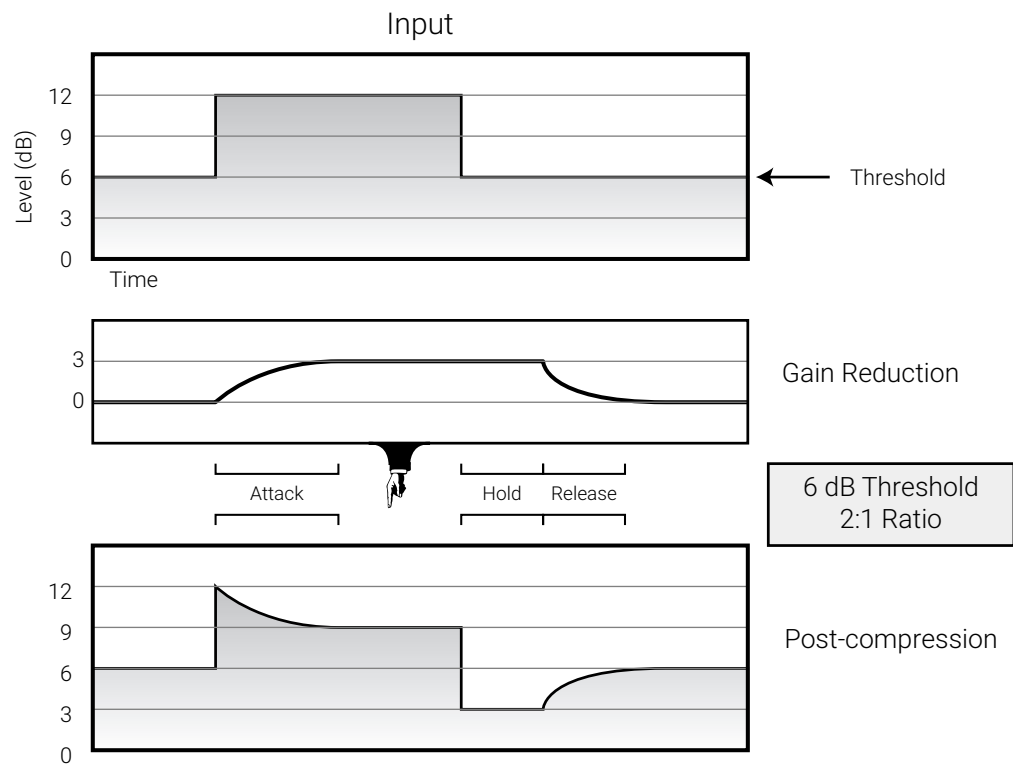
Input

Figure 17.14 The hold function on a compressor. This simplified illustration shows the input levels at the top, the amount of applied gain reduction in the middle, and the post-compression level at the bottom. We can see that once the original level drops, the compressor holds the gain reduction for a while before release.

So, in order to reduce the input level by 3 dB, the gain stage should be fed with –3 dB. In order to convert the gain reduction to gain, the magnitude of the former is mirrored around the 0 dB line using a simple phase inversion (Figure 17.5f).

Additional controls

Now that we have established the basic principles behind the operation of compressors, let's take a look at a few more controls. Figure 17.15 shows the addition of the controls we will look at in the next sections.

Makeup gain

A compressor's principal function is to make louder sounds softer. As a result, the perceived loudness of the compressed signal is likely to drop. To compensate for this, the makeup gain control (sometimes called *gain* or *output*) simply boosts the level of the output signal by a set number of dBs. The boost is applied uniformly to the signal, independently of any other control setting—both signals below and above the threshold are affected. Compressors achieve this by either biasing the side-chain's gain
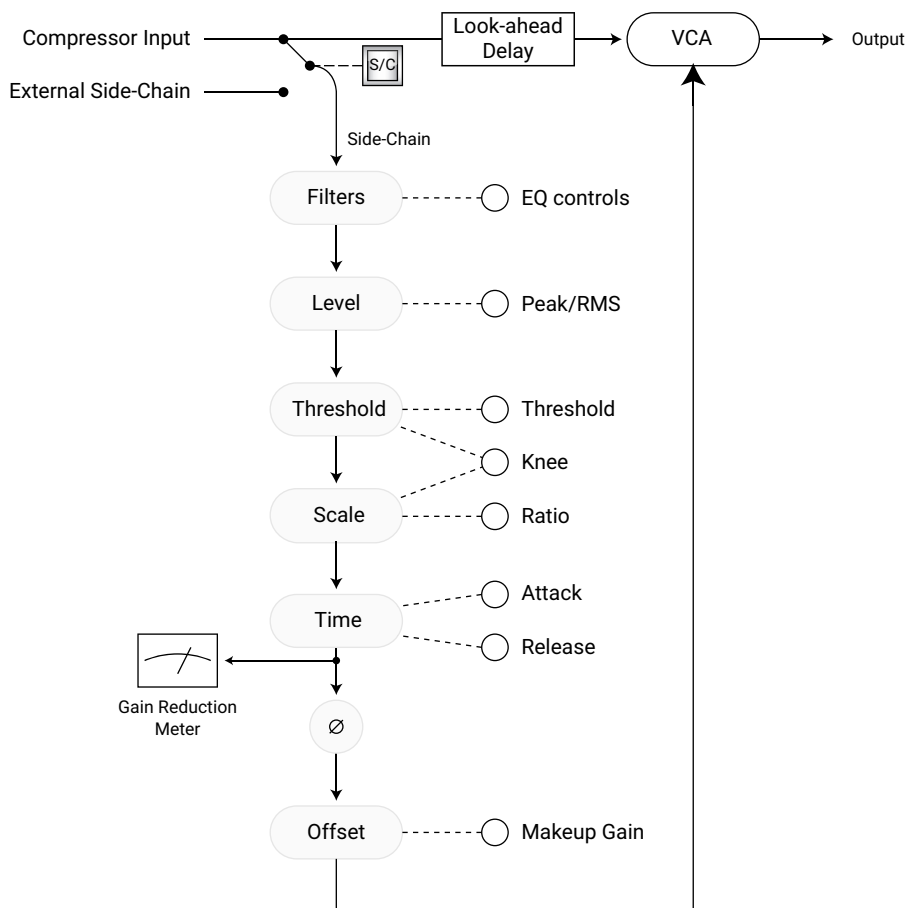
Figure 17.15 Detailed depiction of a compressor's internal workings.

amount before it is applied by the gain stage or by simply amplifying the signal after the gain stage.

As per our louder-perceived-better axiom, when we do A/B comparisons, there is a likelihood that the compressed version will sound less impressive due to the loudness drop. Makeup gain is often set so that, whether the compressor is active or not, the perceived signal loudness remains the same. This way, any comparison made is fair and independent of loudness variations.

Some compressors have an **automatic makeup gain**. Compressors with auto makeup gain calculate the amount of gain required to level the input and output signals based on settings such as threshold, ratio, and release. The auto makeup gain is independent of the input signal and only varies when the compressor controls are adjusted. Arguably, there is no way an automatic makeup gain will match flawlessly the dynamics of all instruments and the possible levels at which these have been recorded. In practice, auto makeup gain often produces a perceived loudness variation when the compressor is bypassed. When able to do so, people will often choose to turn this function off.
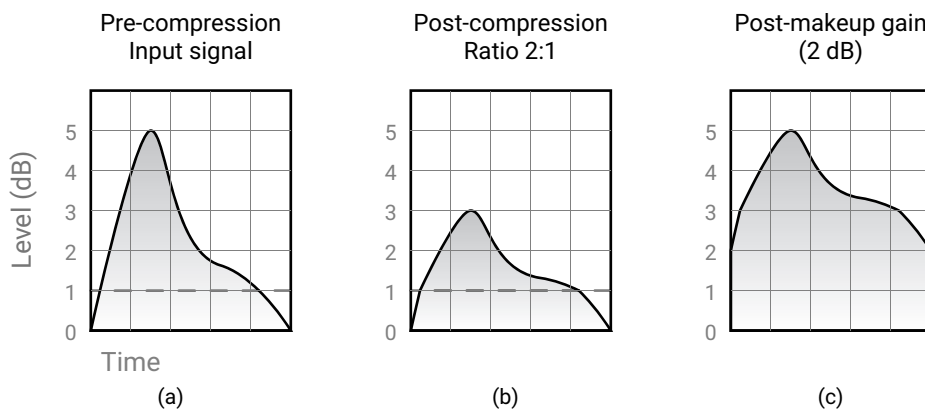
Figure 17.16 Makeup gain. (a) The input signal before compression. (b) The signal after compression with 2:1 ratio, but before makeup gain. (c) The output signal after 2 dB of makeup gain. Note that, in this example, the peak measurement of both the input (a) and the output (c) signals is identical. Yet (c) will be perceived as louder.

Hard- and soft-knees

On the ratio curve, the knee is the threshold-determined point where the ratio changes from unity gain to the set ratio. It takes little imagination to see that it is so named because on a transfer characteristics graph, this curve is reminiscent of a sitting person's knee. The type of compressors we have discussed so far work on the **hard-knee** principle—the threshold sets a strict limit between no treatment and full treatment. The sharp transition between the two provides a more evident compression and a more distinctive effect. We can soften such compression by lengthening the attack and release, but such settings do not always complement the compressed material.

The **soft-knee** principle (also termed *over-easy* or *soft ratio*) enables smoother transition between no treatment and treatment—gain reduction starts somewhere below the threshold with diminutive ratio, and the full compression ratio is reached somewhere above the threshold. While a hard-knee compressor toggles between 1:1 and 4:1 as the signal overshoots, on a soft-knee compressor the ratio gradually grows from 1:1 to 4:1 in a transition region that spreads to both sides of the threshold. This second exception for signals below the threshold being reduced in level is illustrated in Figure 17.17.

Soft-knee is useful when we want more transparent compression (often with vocals). The smooth transition between no treatment and treatment minimizes the compression effect, which in turn lets us dial a higher ratio. Having soft-knee also frees the attack and release from the task of softening the compression effect and lets us dial shorter times on both (which can be useful for applications such as loudening). When we are after the compression effect, a hard-knee is more suitable.

> *Soft-knee is for more transparent compression.*
> *Hard-knee is for more effect.*

Compressors might provide hard-knee only, soft-knee only, or a switch to toggle between the two. Some compressors also provide different degrees of knee rates,
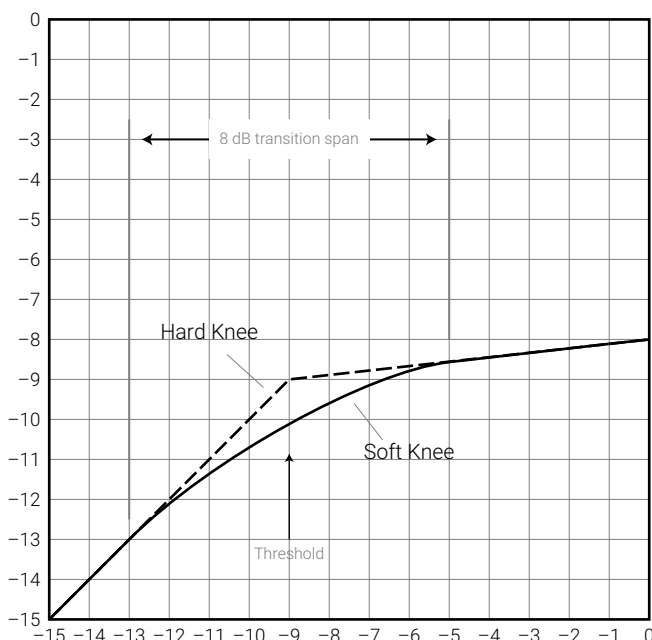
Figure 17.17 Hard- and soft-knees. With hard-knee, the ratio of compression is attained instantly as the signal overshoots. With soft-knee, the full ratio is achieved gradually within a transition region that spreads to both sides of the threshold.

which let us determine the dB span of the transition region with relation to the input scale. For example, in Figure 17.17, the transition span is 8 dB, with 4 dB to each side of the threshold. The compression in this case will start with a gentle ratio at −13 dB, and the full ratio will be achieved at −5 dB. A transition span of 0 dB denotes hard-knee function (such a setting can be seen on the *Oxford Dynamics* plugin in Figure 17.13, below the attack control).

In order to produce soft-knee behavior, an analog compressor has to alter both the threshold and the scale functions. A precise soft-knee is difficult to achieve in the analog domain: the threshold might not fall within the center of the knee, and the ratio is likely to diverge outside the transition region. Digital compressors have no problem in exhibiting a perfect soft-knee like the one in Figure 17.17.

Track 17.16: Vocal Uncompressed
This uncompressed vocal track involves noticeable level variations.

Track 17.17: Vocal Hard-knee
Although the level fluctuations have been reduced, some still exist, mainly due to the operation of the compressor. The working compressor can be easily heard on this track.

Track 17.18: Vocal Soft-knee
Compression settings identical to the previous track, only with a soft-knee (40 dB transition span). The level fluctuations in the previous track are smoothened by the soft-knee, and the operation of the compressor is less evident.

One issue with soft-knee is that the compression starts earlier (below the threshold), and so does the attack. Since the attack starts earlier, less of the natural attack is retained. Therefore, soft-knee might not be appropriate when we try to retain the natural attack of sounds, and with it longer attack times might be needed. Notice the loss of dynamics, life, and some punch in the soft-knee version:

Track 17.19: Drums Hard-knee

Track 17.20: Drums Soft-knee

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack EZdrummer

Look-ahead

Compression can be tricky with sharp level changes, like those of transients. In order to contain transients, a compressor needs a very fast response, but this is not always possible. For one, the gain stage of some compressors, optical ones for instance, is often not fast enough to catch these transients. Then, even if a compressor offers fast response times, the quick clamping down of signals might not produce musical results. It would be great if the side-chain could see the input signal slightly in advance so it could have more time to react to transients. The look-ahead function enables this.

One way to implement look-ahead on analog compressors is to delay the signal before it gets processed. The delay, often around 4 ms long, is introduced after a copy is sent to the side-chain (Figure 17.18). This way, a transient entering the compressor will be seen immediately by the side-chain, but will only be processed shortly after. This enables longer (more musical) attack times since there is a gap of a few milliseconds between the compression onset and the actual processing of the signal that triggered it. By way of analogy, if we could have look-ahead in tennis, it would be like freezing the ball for a while as it crosses over the net, so a player can perfectly reposition after seeing where the ball is heading.

On a hardware compressor, a look-ahead function will be switchable. When look-ahead is engaged, the output signal is delayed as well. These few milliseconds of delay rarely introduce musical timing outflow, but can lead to phase issues if the compressed signal is mixed with a similar track (snare top, snare bottom, for example). With software plugins, any delay that might be introduced is normally auto-compensated.
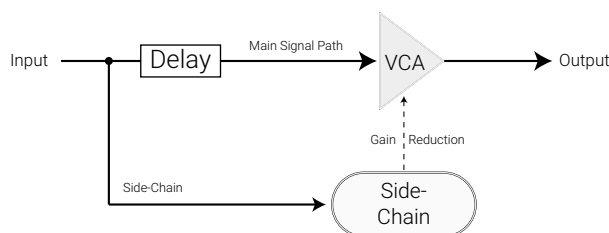
Figure 17.18 A look-ahead function on a hardware compressor. After a copy is sent to the side-chain, the signal is delayed on the main signal path, giving the control circuit more time to respond to the signal.

Stereo linking

A stereo hardware compressor is essentially a unit with two mono compressors. Unless otherwise configured, these two mono compressors work independently in what is known as *dual-mono mode*. Let us consider what happens when stereo overheads are fed into a compressor in such a mode. A floor tom is likely to be louder on the microphone that covers one side of the drum kit—say, the right side. When the tom hits, the right compressor will apply more gain reduction than its left counterpart. As a result, the left channel will be louder and the overheads image will shift to the left. Similarly, when the left crash hits, the overheads image will shift to the right. So, in dual-mono mode, every time a drum hits on either sides of the panorama, the stereo image might shift to the opposite side. Stereo linking interconnects both compressors so an identical amount of gain reduction is applied to both channels. With stereo linking engaged, as the floor tom hits, both sides are compressed to an identical extent and no image shifting occurs.

> *In all but a very few cases, stereo linking is engaged when*
> *a stereo signal is compressed.*

There are various ways to achieve stereo linking. For example, the stereo input might be summed to mono before feeding both the left and right side-chains. The problem with this approach is that phase interaction between the two channels might produce a disrupted mono sum. To combat this, some compressors keep a stereo separation throughout the two side-chains and take the strongest-win approach—the heaviest gain reduction product of either channel is fed to the gain stage of both. With the strongest-win approach, it still makes sense to have different settings on the different channels—by setting a lower ratio on the right channel, we make the compressor less sensitive to right-side events (such as a floor tom hit).

Track 17.21: Drums Stereo Link On
The compression triggered by the tom causes the kick, snare, and hi-hats to drop in level, but their position on the stereo image remains the same.

Track 17.22: Drums Stereo Link Off
With stereo linking off, image shifting occurs. With every tom hit, the kick, snare, and hi-hats shift to the left. As the tom decays, these drums slowly return to their original stereo position.

Plugin: PSP Vintage Warmer 2
Drums: Toontrack EZdrummer

Unlike their hardware counterparts, software compressors rarely provide a stereo linking switch. A software compressor knows whether the input signal is mono or stereo, and, for stereo input, stereo linking is automatically engaged. Some compressors still provide a switch to enable dual-mono mode for stereo signals. Such a mode might be required when the stereo signal has no image characteristics, as in the case of a stereo delay.
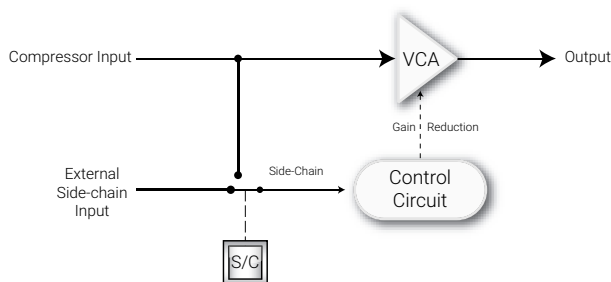
Figure 17.19 External side-chain. An external signal can feed the side-chain, so the compressor input is compressed with relation to an external source.

External side-chain

By default, both the gain stage and the side-chain are fed with the same input signal. Essentially, the signal is compressed in relation to its own level. Most compressors let us feed an external signal into the side-chain, so the input signal is compressed in relation to a different source (Figure 17.19). For example, we can compress a piano in relation to the level of the snare. Hardware compressors have side-chain input sockets on their rear panel, while software sequencers let us choose an external source (often a bus) via a selection box on the plugin window. A switch is often provided to toggle between the external side-chain and the native, internal one. There is also usually a switch that lets us listen to the side-chain signal via the main compressor output (overriding momentarily the compressed output).

Side-chain filters and inserts

Quite a few compressors let us equalize the side-chain signal (the section to the right of the plot in Figure 17.4 is one example). As can be seen from Figure 17.5, such equalization occurs on the side-chain only (not the main signal path) and affects both the native and external side-chain sources. Some analog compressors also provide a side-chain insertion point, enabling processing by external devices. Soon we will look at the many applications for side-chain equalization.

Compressor meters

Three metering options are commonly provided by compressors: input, output, and gain reduction. When it comes to mixing, people often say, "Listen, don't look." True—we should be able to set all the compressor controls based on what we hear; but recognizing the subtleties in compressor action takes experience, and the various meters can often help, especially when it comes to initial rough settings. Fine and final adjustments are done by ear.

Out of the three meters, the most useful is the **gain reduction meter**, which shows the applied amount of gain reduction. The gain reduction reading combines the effect of all the compressor controls; for example threshold, ratio, attack, release, peak, RMS, etc. The main roles of this meter are:

• To teach us **how much** gain reduction is applied (which hints at the most suitable amount of makeup gain).

- To teach us **when** gain reduction takes place. Mostly we are interested in:
  - when it starts; and
  - when it stops.
- To provide a visual indication of the **attack** and **release** activity (which can help in adjusting them).

To give an example, when our aim is to balance the level of vocals, we want the gain reduction meter to move with respect to the level fluctuation that we hear on the unprocessed voice (or see on the input meter). If the attack or release is too long, the meter will appear lazy compared with the dynamics of the signal, and so might the compression.

The terms *compression* and *gain reduction* are often used interchangeably, where people think of the amount of compression as the amount of gain reduction. When people say they compress something by 8 dB, they mean that the gain reduction meter's highest reading is 8 dB. Associating the gain reduction reading with the amount of compression is not a good practice, though. If the gain reduction meter reads steady at 8 dB (say, due to a hypothetical release time of 10 hours), no compression is taking place—there is simply constant gain attenuation of 8 dB. For the most part, such a compressor behaves like a fader—apart from when the gain reduction climbs to 8 dB, the output signal keeps all its input level variations, whether these happened above or below the threshold. In order for compression to occur, the amount of gain reduction must vary over time and the gain reduction meter must move; and the faster it moves, the more compression takes place. There is a similarity here to sound itself—sound is the outcome of changes in air pressure; steady pressure, whether high or low, does not generate sound. Likewise, the amount of compression is determined by changes in gain reduction. We might want to consider the amount of compression as an average of the absolute differ- ence between the peak and RMS gain reduction. But that's a story for a different book.

In addition, we should also consider the range within which this meter moves. If a vocal overshoots the threshold between phrases, but within each phrase the gain reduction meter only moves between 4 and 6 dB, then the effective compression is only 2 dB. At the beginning and end of each phrase, the signal is still compressed on its way to and back from 4 dB, but this compression is marginal compared with the compression taking place between 4 and 6 dB. We will call the range where gain reduction varies *dynamic compression* (4–6 dB in our case), and use the term *static compression* for the static range (0–4 dB in our case), where gain reduction only occurs on the way to or back from the dynamic compression. Figure 17.20 illustrates the differences between the two. Static compression can happen either as a result of slow release or a threshold being set too low. If the latter is the case, it might be wise to bring up the threshold, so the gain reduction meter constantly moves between 0 dB and the highest reading. This would minimize static compression. However, there are cases where static compression makes perfect sense, such as when compression is employed to shape the dynamic en- velope of sounds and the threshold is set intentionally low to enable more apparent attack.
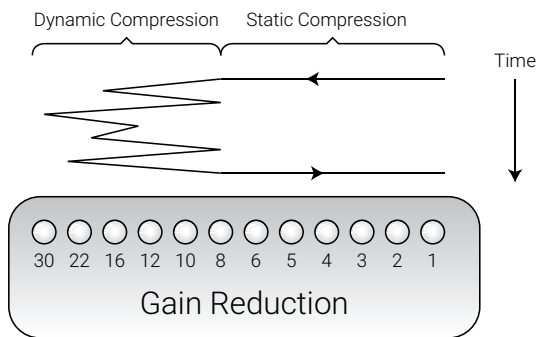
Figure 17.20 Dynamic and static compression. The dynamic compression happens in the range where the gain reduction meter varies. The static compression only happens as gain reduction rises to or falls from the dynamic compression range.
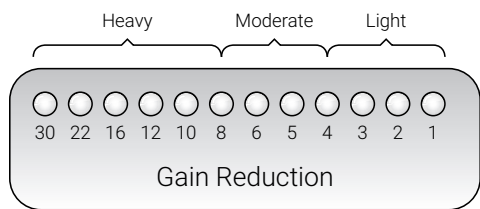
Figure 17.21 A gain reduction meter. The descriptions above the meter are an extremely rough guideline as to what some people might associate different amounts of gain reduction with.

The subjective observation of light, moderate, or heavy compression is based on the gain reduction meter. We just saw that the amount of compression is far more complex than the basic reading this meter offers. It is also unsafe to translate any observations to numbers, as what one considers heavy the other considers moderate. We can generalize that 8 dB of compression on the stereo mix is quite heavy. For instruments, people usually think of light compression as somewhere below 4 dB, heavy compression as somewhere above 8 dB, and moderate compression as somewhere in between these two figures (Figure 17.21). We are talking about dynamic compression here—having 20 dB of constant gain reduction is like having no compression at all.

You might have noticed that the numbers in Figure 17.21 are positive and arranged from right to left. This is the opposite direction to the input and output meter, and suggests that gain reduction brings the signal level down, not up. Also, on compressors with no dedicated gain reduction meter, gain reduction is displayed on the same meter as the input and output levels (with a switch to determine which of the three is shown). Since on most meters the majority of the scale is below 0 dB, this is also where gain reduction is shown. Very often a gain reduction of 6 dB will be shown as –6 dB.

To be technically correct, a decrease of a negative value is an increase, so a gain reduction of –6 dB essentially denotes a gain boost of 6 dB. Many designers overlook that, and, despite having a dedicated meter labeled "gain reduction," the scale incorrectly shows negative values.

The **input meter** can be useful when we set the threshold. For example, if no signal exceeds –6 dB except for a peak we want to contain, this is where the threshold might go. We can also determine the action range of input levels by looking at this meter (more on action range below). **The output meter** gives a visual indication of what we hear and ensures that the output level does not exceed a specific limit, such as 0 dBFS. The output meter can also help us in determining rough makeup gain without having to toggle the compressor in and out.

Controls in practice

Threshold

Figure 17.22 shows possible levels of a non-percussive performance—a vocal phrase, perhaps. We can see that the levels fluctuate between the 0 and –10 dB range. We will call this range the *action range* and it can help us determine the threshold setting for specific applications. It should be made clear that setting the threshold anywhere above 0 dB has little point—unless the knee is soft, no compression will take place. Setting the threshold within this range will yield selective compression—only the signal portions above the threshold will be treated. Sometimes this is what we want, as when we only want to treat loud parts of the signal. However, the danger in selective compression is that the transition between treatment and no treatment happens many times, and the compression becomes more obvious. Setting the threshold to the bottom of the action range (–10 dB in Figure 17.22) will result in uniform treatment for all the level changes within it. While this might not be suitable with every compressor application, when a performance involves an evident action range, its bottom might provide a good starting point for threshold settings.

One key point to consider is what happens when we set the threshold below the action range. Say we first set the threshold to the bottom of the action range at –10 dB, and set the ratio to 2:1. The highest peak in Figure 17.22 hits 0 dB, so by overshooting 10 dB it will be reduced by 5 dB. Also, if we assume that the compressor timing rates are based on 10 dB and the release is set to 10 ms, as the same peak dives to –10 dB, it will take the 5 dB of gain reduction 5 ms to recover. Now if we reduce the threshold to –20 dB, the same peak overshoots by 20 dB and will be reduced by 10 dB—twice as much as before. Also, as the gain reduction is now 10 dB, it will take it 10 ms to recover. Effectively, by lowering the threshold we increase the amount of gain reduction, which results in further depression of dynamics. In addition, we slow down
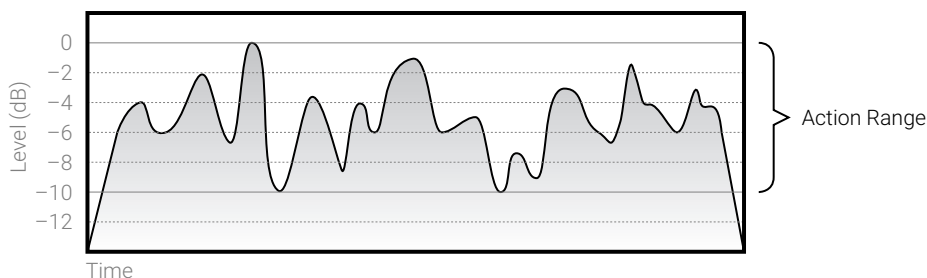
Figure 17.22 The action range of a non-percussive performance. Apart from the initial rise and final drop, the signal level is only changing within the 0 to –10 dB range.
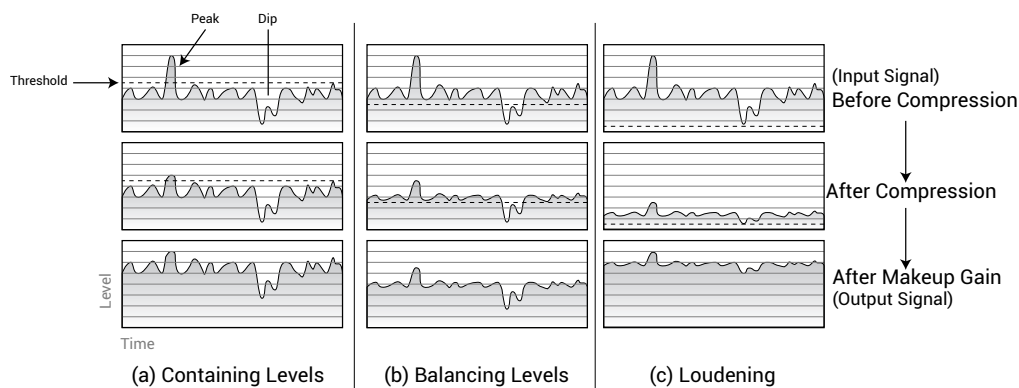
Figure 17.23 Threshold settings for common compression applications. (a) Containing levels. (b) Balancing levels. (c) Loudening. The columns from top to bottom show the input signal and the threshold setting, then the compressed signal before makeup gain and the output signal after makeup gain.

the release (and attack) response. Also, having the threshold at –20 dB will result in 5 dB of static compression as everything between –20 and –10 dB will be constantly compressed (apart from during the leading rise and closing fall).

Let us discuss three common compression applications and their threshold settings. These are containing peaks, level balancing, and loudening (condensing levels). Figure 17.23 shows the level of a hypothetical performance, say vocal. Each column represents a different application, and the threshold setting for each is shown on the top row. The top row also shows the input signal, which is identical in all three cases. In this specific performance, there is one peak and one level dip.

The left column (a) is concerned with containing levels. The high threshold is set around the peak base, above all other levels. The compression reduces the peak level, but no other portion of the signal. Reducing the peak lets us boost the overall level of the signal, and, by matching the original peak with the compressed one, moderate- and low-level signals are made louder. Note that in this case, the loss of dynamics is marginal.

The term containing levels *is being used here as a more general term for* containing peaks*. Containing peaks is concerned with preventing levels from exceeding a predefined limit—a job for a limiter, really. We might want to contain the louder downbeats of a strumming guitar, although these might not peak.*

The center column (b) demonstrates possible compression for level balancing. The threshold is set to a moderate level right above the dip, but below all other portions of the signal. The idea is to pull everything above the dip down toward it, which is exactly what happened after compression. Note, however, that compared with the compression shown on the left column, the signal has lost more of its dynamics. The makeup gain offsets the signal, so moderate levels are back to their original area (around the middle line of the input graph). If we look at the output graph, we can see that the peak was pulled down toward this middle line, while the dip had a gentle push up from the
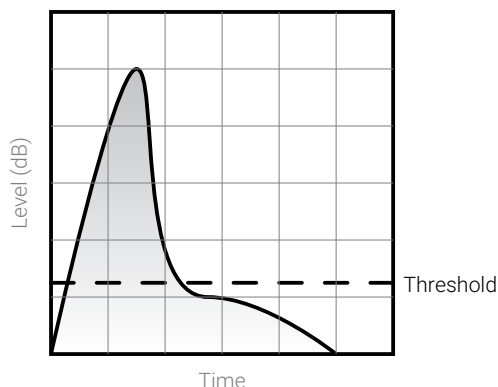
Figure 17.24 Catching the attack of a snare. The threshold is set so the attack is above it but the decay is below.

makeup gain. Again, one possible issue with this technique is the uneven compression effect it deposits due to selective compression—the dip gets louder, but as opposed to the rest of the signal its dynamics remain unaffected. If the applied compression results in audible effect, it will not be evident on the dip. To overcome this, we can lower the threshold so even the dip gets a taste of the compressor effect.

The right column (c) exhibits a compression technique used when we wish to make instruments louder or when we want to condense their dynamics. The low threshold is set at the base of the signal, so all but the very quiet levels get compressed. Two things are evident from the post-compression graph. First, the levels are balanced to a larger extent as the dip was treated as well. After compression, the peak, dip, and moderate levels are highly condensed. Thus, this technique can also be used for more aggressive level balancing. Second, we can see that the signal has dropped in level substantially. To compensate for this, makeup gain is applied to match the input and output peaks. In this case, the output signal will be perceived as louder than the input signal, since its average level rose by a fair amount. The main issue with this technique is that, of all three applications, it led to the greatest loss of dynamics.

There is more to threshold settings than these general applications. Sometimes we know which part of an instrument's dynamic envelope we should treat. On a snare, for example, it might be the attack, so the threshold is set to capture the attack only. Figure 17.24 illustrates this. One important thing to note is that the lower the threshold, the sooner we catch the attack buildup. In turn, this lets us lengthen the attack time, which, as we shall soon see, is often beneficial.

Ratio

A compressor with ratios above 10:1 behaves much like a limiter—with a 10:1 ratio, a signal overshooting by a significant 40 dB is reduced to a mere 4 dB above the threshold. Yet, even a compressor with an ∞:1 ratio is not really a limiter (signals can still overshoot due to long attack or the RMS function), so ratios higher than 10:1 have a place in mixes. Mastering ratios (on a compressor before the limiter) are often gentle and kept around 2:1. Mixing ratios can be anything—1.1:1 or 60:1. Yet the ratio 4:1 is often quoted as a good starting point in most scenarios.

Logic has it that the higher the ratio, the more the compression applied and the more obvious its effect. One important thing to understand is that the degree of compression diminishes as we increase the ratio. As can be seen in Figure 17.25, turning the ratio from 1:1 to 2:1 for an 8 dB overshoot results in gain reduction of 4 dB. Turning the ratio from 8:1 to 16:1 for the same overshoot only results in an additional gain reduction of 0.5 dB. A more practical way to look at this is that changes of lower ratios are much more significant than the same changes for higher ratios—a ratio change from 1.4:1 to 1.6:1 will yield more compression than a ratio change from 8:1 to 16:1. Most ratio controls take this into account by having an exponential scale (for example, 1:1 to 2:1 might be 50 percent of the scale; 2:1 to 4:1 will be an additional 25 percent, and so forth).
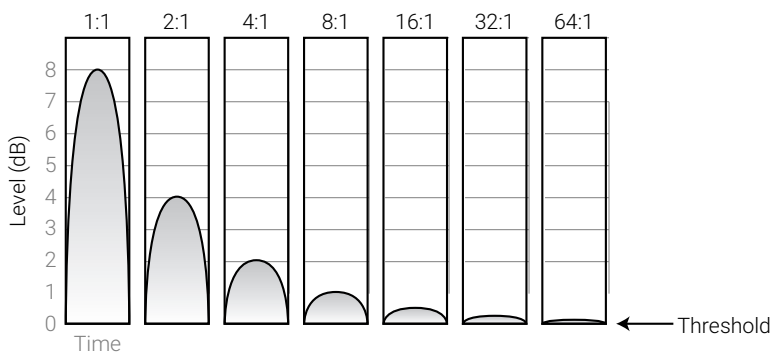
Figure 17.25 The diminishing effect of increasing the ratio. The ratio is shown above each graph. The threshold is set to 0 dB.

A demonstration of the diminishing effect of increasing the ratio can be heard in Track 17.3 (Rising Ratio). The majority of level lost and the increasing effect of compression is mostly evident in the first bar, where the ratio rises from 1:1 to 2:1, and in the second bar, where the ratio further rises to 4:1. The change in the following bars is marginal compared to the change in these first two bars.

We can make some assumptions as to the possible ratios in the three applications discussed in Figure 17.23. In the first case—containing levels—a high ratio, say 10:1, might be suitable. With the high threshold set to capture the peak only, the ratio simply determines the extent by which the peak is brought down, with a higher ratio being more forceful. In the second case—balancing levels—the moderate threshold is set around the average level. The function of the ratio here is to determine the balancing extent. Since most of the signal is above the threshold, a high ratio would result in a noticeable effect and possible oppression of dynamics. However, a low ratio might not give sufficient balancing, so a moderate threshold, say 3:1, could be suitable here. In the third case—loudening—the low threshold already sets the scene for some aggressive compression. Even a moderate ratio can make the compression so aggressive that it will suppress dynamics. Therefore, a low ratio, say 1.4:1, might be appropriate.

The relationship between threshold and ratio

In the three examples above, the lower the threshold, the lower the ratio, which brings us to an important link between the two controls. Lowering the threshold means more compression, while lowering the ratio has the opposite effect. Often these two controls are fine-tuned simultaneously, where lowering one is followed by the lowering of the other. Ditto for rising. The idea is that once the rough amount of required compression is achieved, lowering the threshold (more compression) and then lowering the ratio (less compression) will result in roughly the same amount of compression, but with a slightly different character.

Although a lower threshold and higher ratio each result in more compression, the effect of each is different. A lower threshold results in more compression as larger portions of the signal are compressed. A high ratio results in more compression on what has already been compressed. Put another way, a lower threshold means more is affected, while a higher ratio means more effect on the affected. If the threshold is set too high on a vocal take, even a ratio of 1,000:1 will not balance out the performance. If the ratio is set too low, the performance will not balance, no matter how low the threshold is. Finding the right balance between threshold and ratio is one of the key nuts to crack when compressing.

> *Threshold defines the extent of compression. Ratio defines the degree.*
> *The two are often fine-tuned simultaneously.*

This might be the right point to talk about how the two are utilized with relation to how evident the compression is. We already established that the threshold is a form of discrimination between what is being treated and what is not. In this respect, the higher the threshold, the more the discrimination. We can consider, for example, the compression of vocals, when our task is to make it as transparent as possible. Chances are that, with moderate threshold and moderate ratio, the effect of compression would be more evident than if we set the threshold very low but applied a very gentle ratio: 1.2:1, for example (Figure 17.26). Since we already established that the attack and release are also applied on level changes above the threshold, we know that the timing characteristics are applied even if the threshold is set very low. All of this is said just to stress that a lower threshold does not necessarily mean more evident effect—this is determined by both the threshold and the ratio settings, as well as other settings described below.

Attack

One reason for longer attack times was given earlier—to prevent the suppression of an instrument's natural attack. But we don't always want to keep the natural attack. Sometimes we want to soften a drum hit so kicks, snares, and the like are less dominant in the mix. In such cases, shorter attack is dialed.

> *The longer the attack time, the more we retain the instrument's*
> *natural attack. However, this is not always wanted.*

The way longer attack helps to retain the natural attack is easily demonstrated on percussive instruments. Figure 17.27 illustrates the effect of different attack times on a snare. We can see that, with no attack, the natural attack is nearly lost as the compressor
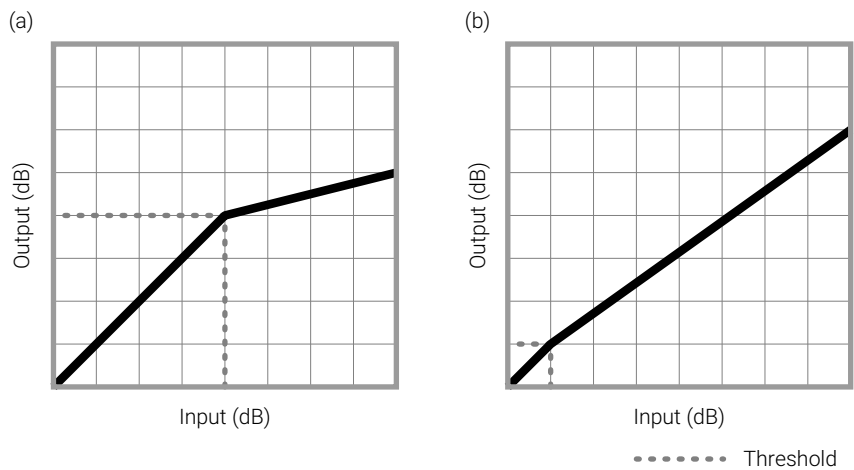
Figure 17.26 The effect of compression with relation to the threshold and ratio settings. The scale in these graphs represents the full dynamic range of the input signal. (a) Moderate threshold and ratio. One should expect more evident compression here as the threshold sets a discrimination point between untreated signals and moderately treated signals. (b) Low threshold and ratio. One should expect less evident compression here as all but the very quiet signal will be treated gently.
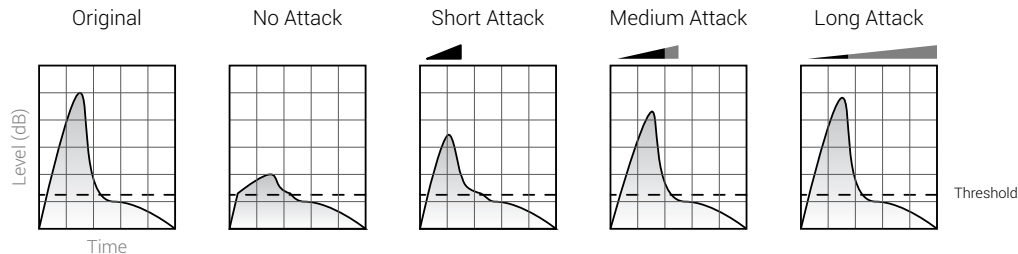
Figure 17.27 Different attack times on a snare. The attack buildup is shown below each caption. The longer the compressor attack time, the more of the instrument's natural attack is retained.

squashes it down quite drastically (clearly, this is also dependent on the threshold and ratio settings). Then, the longer the attack, the more of the natural attack is retained. A long attack hardly affects the natural attack—or, in this specific snare case, hardly affects anything at all.

We might instinctively think: If we want to retain the natural attack of an instrument, why not just set the attack time as high as possible? The long attack in Figure 17.27 demonstrates why this is unwise. We can see that, by the time the signal has dropped below the threshold, the attack has barely built up to cause any significant gain reduction—it nearly canceled out the compression effect, making the whole process rather pointless. Based on Figure 17.12, we know that, shortly after the snare signal starts to drop, the applied gain reduction also starts to drop, and at this point the attack phase stops and the compressor enters the release phase. This point on Figure 17.27 is where the attack buildup turns gray, and we can see that the potential full effect is not achieved for either the medium- or the long-attack settings.

Track 17.23: Snare Uncompressed
The original snare track used for the following samples. For the purpose of demonstration, the compressor threshold in the following tracks was set to –15 dB and the ratio to ∞:1.

Track 17.24: Snare Attack 10 ms
The extremely short attack on the compressor immediately reduces the natural attack of the snare, simply resulting in overall level reduction.

Track 17.25: Snare Attack 1 ms
With 1 ms of attack, the result is not much different from in the previous track. Yet, a very small part of the natural attack manages to remain, so there is a notch more attack on this track.

Track 17.26: Snare Attack 10 ms
Ten milliseconds is enough for a noticeable part of the natural attack to pass through, resulting in compression that attenuates the late attack and the early decay envelopes.

Track 17.27: Snare Attack 50 ms
An even larger portion of the natural attack passes through here, resulting in gain reduction that mainly affects the decay.

Track 17.28: Snare Attack 100 ms
One hundred milliseconds of attack lets the full natural attack through. Essentially, little gain reduction is applied and it affects the early decay envelope. Compared to the uncompressed track, a bit of the early decay was reduced in level here.

Track 17.29: Snare Attack 1,000 ms
The majority of sound in each of these snare hits is no longer than 500 ms. One second is too long for the compressor to respond, so no gain reduction has been applied at all, resulting in a track that is similar to the uncompressed one.

Plugin: PSP MasterComp
Drums: Toontrack EZdrummer

Even when the signal remains above the threshold after the attack phase, a longer attack time is rarely beneficial. We can demonstrate this on the piano key hit in Figure 17.28. The short attack brings the gain reduction to full impact slightly after the natural attack has declined, and the resultant dynamic envelope resembles the input signal envelope. This is not the case with the long attack, which clearly altered the dynamic envelope and resulted in unwanted timbre alteration. This example teaches us what is very likely to be a musical attack time for instruments of a percussive nature: one that only affects the natural attack without altering the dynamic envelope succeeding it.

Another reason why a longer attack would be beneficial involves the compression of overheads. If the kick is the main instrument to trigger the compression and the attack is set too short, not only the kick's natural attack but also any other drum played underneath it (such as the hi-hats) would be suppressed. If this happens on the downbeat, downbeat hats would be softer than other hats. This might be suitable with more exotic rhythms, but rarely for the contemporary acid-house track, punk-rock,
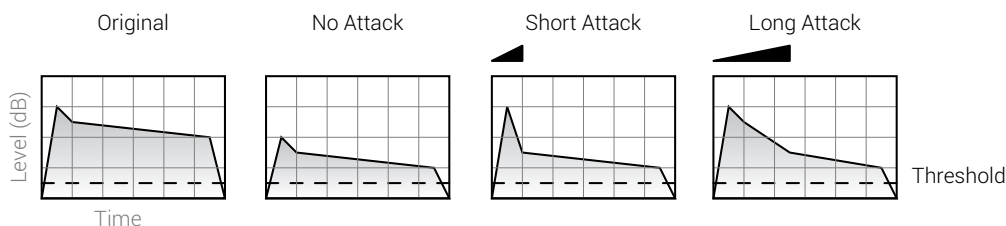
Figure 17.28 Different attack times on a piano key hit. The long attack time caused alteration of the level envelope, which would distort the timbre of the instrument.

gangsta rap, or just about anything else. A longer attack will let the downbeat hats sneak in before the compressor acts upon them.

One of the things we have to consider when it comes to attack times is how much gain reduction actually takes place. The more gain reduction there is, the more noticeable the attack will be—there is no comparison between an attack on maximum gain reduction of 1 dB and that of 8 dB. Generally speaking, heavy gain reduction could benefit from longer attack times, so large upswings take place gradually rather than promptly, making level changes less drastic. In addition, fast attack and serious gain reduction can easily produce an **audible click**. Although such attack clicks are sometimes used to add definition to kicks, they are mostly unwanted.

Nevertheless, there is a catch-22 here since a longer attack lets some of the level through before bringing it down (and with heavy compression it is quite a long way down). So with longer attack times level changes might be less drastic but more noticeable. This is often a problem when heavy compression is applied on less percussive sounds such as vocals.

Track 17.30: Drums Uncompressed
The uncompressed version of the drums used in the following samples.

Track 17.31: Drums Attack 5 ms
Notice how the hi-hats' level is affected by the compression applied on the kick and snare. Also, the short attack suppresses the natural attack of both the kick and the snare.

Track 17.32: Drums Attack 16 ms
Sixteen milliseconds of attack is enough to prevent the kick and snare compression from attenuating the hats. In addition, the natural attack of both these drums is better retained.

Plugin: McDSP Compressor Bank CB1
Drums: Toontrack EZdrummer

Track 17.33: Vocal Attack 1 ms
This very short attack time yields instant gain reduction, which makes the compression highly noticeable. Essentially, each time the voice shoots up, a quick level drop can be discerned.

Track 17.34: Vocal Attack 5 ms
Five milliseconds of attack means that the gain reduction is applied more gradually, producing far fewer level fluctuations, due to the operation of the compressor. This attack time is also short enough to reasonably tame the vocal upswings.

Track 17.35: Vocal Attack 7 ms
There are still some quick gain reduction traces in the previous track, and 7 ms of attack reduces these further. However, the longer attack means that the voice in this track manages to overshoot higher.

Track 17.36: Vocal Attack 30 ms
Thirty milliseconds is too long, and the compressor misses the leading edge of the vocal upswings. In addition, the compressor does not track the level variations of the vocal, making level fluctuation caused by the compressor more noticeable.

Track 17.37: Vocal Attack Pop
This track is also produced using 30 ms of attack, but with a lower threshold and higher ratio so as to draw heavier gain reduction. Profound level pops can be heard on "who," "running," and "pass."

Plugin: Sonnox Oxford Dynamics

If the attack is too short, the vocal dynamics flatten, but if the attack is made longer, there might be a **level pop** on the leading edge of each phrase. It is sometimes possible to find a compromise in these situations, and, when suitable, raising the threshold or lowering the ratio can also help. Yet, a few compressor tricks described later in this chapter can give a much more elegant solution and better results.

Another consequence of a fast attack is **low-frequency distortion**, the reason being that the period of low frequencies is long enough for the compressor to act within each cycle rather than on the overall dynamic envelope of the signal. Figure 17.29 shows the outcome of this, and we can see the attack affecting every half a cycle. The
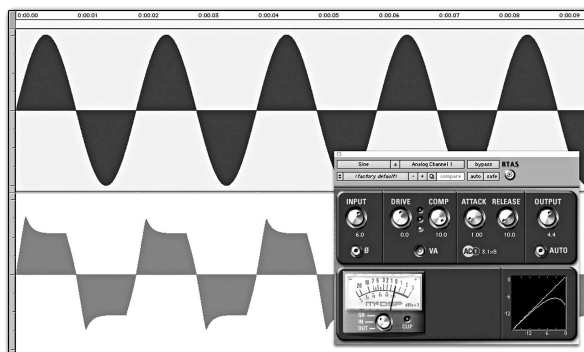
Figure 17.29 Low-frequency distortion resulting from fast attack settings. The top track is a 50 Hz sine wave passing through the McDSP *Analog Channel 1* shown on the bottom right. The compressor setting involves 1 ms of attack and 10 ms of release. The rendered result is shown on the bottom track, where we can see the attack acting within each half cycle.

compressor has reshaped the input sine wave into something quite different and, as a result, distortion is added. The character of such distortion varies from one compressor to another, and it can be argued that analog compressors tend to generate more appealing distortion. In large amounts, this type of distortion adds rasp to the mix. But in smaller, more sensible amounts, it can add warmth and definition to low-frequency instruments, notably bass instruments.

The main purpose of the hold control is to rectify this type of distortion. By holding back the gain reduction for a short period, the attack and release are prevented from tracking the waveform. The cycle of a 50 Hz sine wave is 20 ms; often the hold time is set around this value.

The attack settings can affect the tonality of the compressed signal, where longer attack times tend to soften high frequencies. Alternatively, this can be seen as low-frequency emphasis and added warmth. The reason for this is demonstrated in Figure 17.30. High frequencies are the result of quick level changes, and low frequencies are the result of slow ones. A long release can slow the rate at which the signal level rises, which softens the quick level rise of high frequencies.
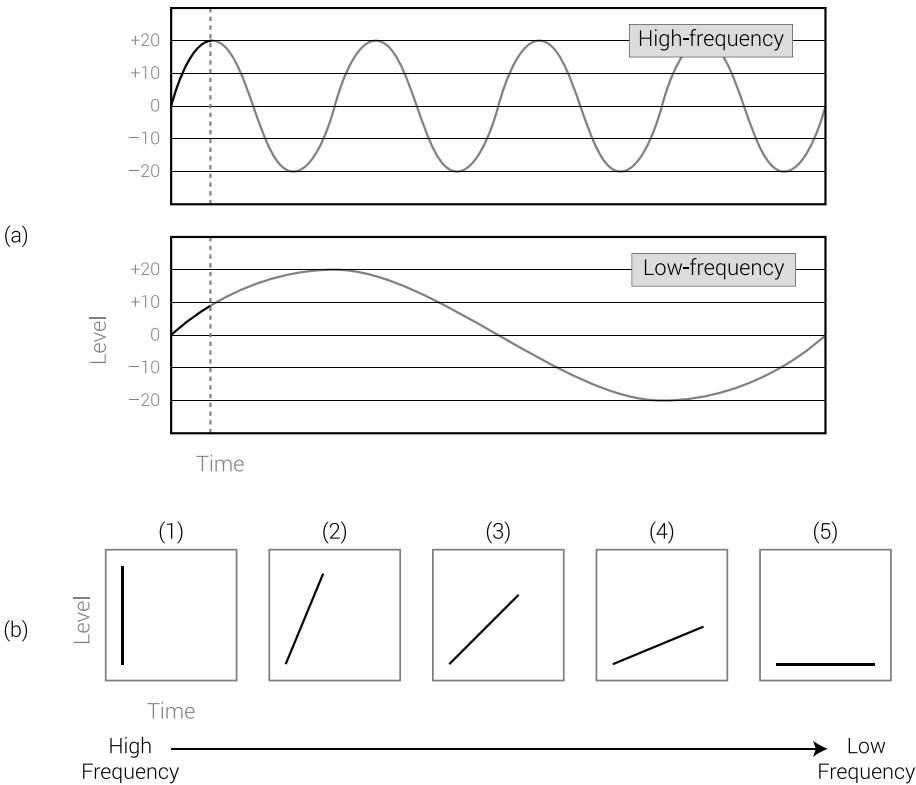
Figure 17.30 Level change rate and frequencies. (a) We can see that the slope of the high-frequency waveform is steeper. By the time the high-frequency sine wave had risen by 20 dB, the low-frequency sine wave had only risen by 10 dB. (b) Faster level changes are associated with high frequencies. A long attack on a compressor can turn a level rise like (2) into (3), which softens high frequencies. We will see later that a gate can result in an abrupt level jump like (1), which generates high frequencies.

Track 17.38: Bass Uncompressed
The source bass track used in the following tracks.

Track 17.39: Bass LF Distortion Fast TC
The low-frequency distortion was caused by the fast attack (10 μs) and fast release (5 ms).

Track 17.40: Bass LF Distortion Fast Attack
Although the release is lengthened to 100 ms, the 10 μs attack still produces some distortion. When a bass is mixed with other instruments, a degree of distortion as such can add definition and edge.

Track 17.41: Bass LF Distortion Slow TC
With the release at 100 ms and the attack lengthened to 2 ms, there is only a faint hint of distortion.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3

Release

Release and attack have a few things in common. A very short release can also **distort low frequencies** for the same reason a short attack can. Here again, the hold control can reduce this type of distortion. A long release tends to soften high frequencies on the same principle that long attack does. A very short release can cause **audible clicks**, although, in contrast to attack clicks, it is hard to think of any practical use for these release clicks. Release settings also become somewhat more important with heavier compression, as the release is applied on larger amounts of gain reduction.

Short release times can cause low-frequency distortion and audible clicks.

Track 17.42: Bass LF Distortion Fast Release
The distortion on this track is the result of a fast release of 5 ms and a relatively long attack of 2 ms.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3

One thing we have to consider is that the gain recovery during the release phase sometimes happens during silent sections. For example, if you try to sing aloud "when we eat" and then "when we go," it is likely that you will cut the "eat" while fading out the "go." Quick vocal drops such as "eat" can be shorter than 50 ms, and a release longer than that period might end up affecting silence. Although the release setting is important to control vocal drops, we have to acknowledge that some of the gain recovery can happen on inaudible sections. This suggests that the release can be more important for instruments that rarely dip into silence, such as overheads and rhythm guitars. These types of instruments are also more likely to experience **pumping**, which is caused by short release times.

Short release can cause pumping.

As per the discussion of the difference between gain reduction and compression, we know that a very long release time (hypothetically, 10 hours) will cause steady gain reduction, which is exactly the opposite of compression. Both long attack and long release slow down changes in gain reduction, which makes a compressor more of a fader. Subsequently, the effect of compression becomes less noticeable. By way of analogy, if during manual gain-riding the engineer's fader movements are extremely slow, the input level variations will hardly alter. The true effect of compression only happens when the gain reduction tracks change in input levels. Short release times result in faster gain recovery, which in turn increases the average output level and the perceived loudness.

*Short release times result in more compression,
more evident effect and loudness increase.*

Two compressed versions of Track 17.30. Note that the compression is more evident in the short-release version, which might also be perceived as louder. With 2,000 ms of release, most of the gain reduction is static, and the result is very similar to the uncompressed version.

Track 17.43: Drums Release 50 ms

Track 17.44: Drums Release 2,000 ms

Plugin: Focusrite Liquid Mix (Trany 1a emulation)
Drums: Toontrack EZdrummer

The previous attack section started with a snare demonstration, so let us start by using a snare to demonstrate the function of release. We established that the longer the attack is, the more the natural attack is retained. The effect of release can easily be misunderstood as it is exactly opposite in nature—the longer the release, the less natural decay is retained. The release control determines the rate at which the gain reduction can drop, or, in other words, how quickly the gain recovers. Once the input signal level drops, a longer release will slow down the recovery rate, so more gain reduction will be applied on signals. This effect was demonstrated in Figure 17.11 on a noise burst; Figure 17.31 shows this happening on a snare. It is interesting to note the effect of a medium release, which quite noticeably reshapes the decay. This extra bump could alter the timbre of a snare in an unflattering way. So often we seek to either leave the decay as is or attenuate it altogether.

The following tracks demonstrate the effect of different release times. In all of them, 50 ms of attack allows most of the natural attack to pass through, leaving the gain reduction to mainly affect the decay of the snare. The longer the release, the slower the gain recovers, and consequently the more the decay is attenuated.

Track 17.45: Snare Release 100 ms

Track 17.46: Snare Release 500 ms

Track 17.47: Snare Release 1,000 ms

Track 17.48: Snare Release 2,000 ms

Track 17.49: Snare Release 3,000 ms
Note that in Track 17.46, it is possible to hear the decay recovering after being reduced in level straight after the attack. This is a demonstration of decay reshaping similar to that caused by the medium release shown in Figure 17.31.

Plugin: PSP MasterComp
Drums: Toontrack EZdrummer
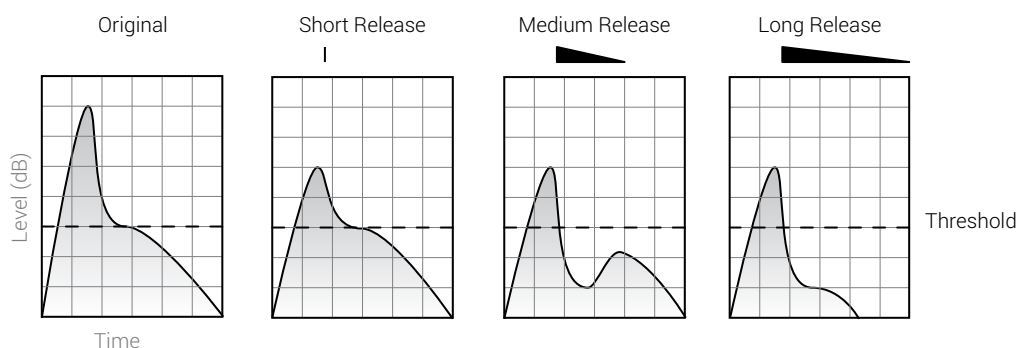
Figure 17.31 Release times on a snare. With a short release, the decay is unaltered compared to the original signal. With a medium release, it takes some time for the gain to recover, and only then the natural decay returns to its original levels. With a long release, the decay is attenuated and very little of it is retained.

In the case of vocals, we might want a fast release so that the already dropping signal is not attenuated further by the compressor. But both attack and release are also concerned with level changes above the threshold, so in the case of vocals their function is often more concerned with tracking level fluctuations. As already suggested, the two are set to be fast enough to track these fluctuations.

Longer release on a compressor means less of the instrument's natural decay and can cause timbre alterations.

No discussion about the release control would be complete without mentioning its strong link to the rhythmical nature of the song. When compression is applied on instruments with a percussive nature, the release dictates the length of each hit and how it decays. Different release settings can make a world of difference to how an instrument locks to the natural rhythm of the music. For example, a snare that hits every half-note might feel right in the mix if its perceived length is set to a quarter-note. There are no hard and fast rules—we have to experiment with the release settings until it feels right.
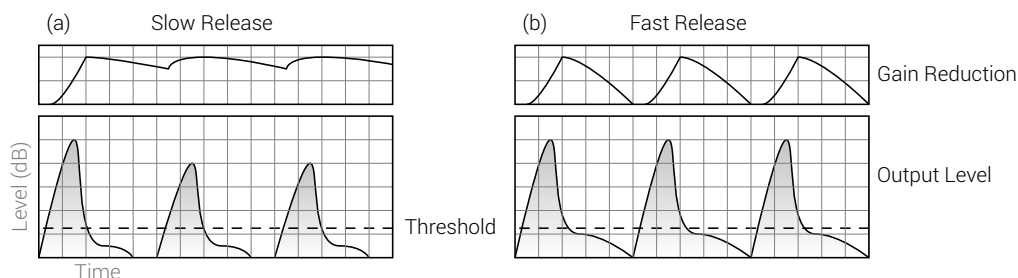
Figure 17.32 Release and snare spacing. (a) With a slow release, the gain reduction does not recover before the next snare hit, causing attenuation to the consecutive hits and softening their attack. (b) With a fast release, the gain reduction is back to 0 dB before the next hit, so the compressor treats all hits uniformly.

After a snare hit, the gain reduction meter shows when (and if) gain reduction falls back to 0 dB. Usually, we want this to happen before the next snare comes along, or the next hit will be attenuated by the ongoing gain reduction (Figure 17.32). One would naturally ask: But what happens if the drummer plays half-notes for the most part, and just before each chorus plays eighth-notes? A release setting that fits the length of a half-note will be too long for eighth-notes and will de-emphasize them. Chances are that these eighth-notes are there to squeeze some excitement in before the chorus, so such a de-emphasis is contrary to what we want. However, setting the release to fit eighth-notes might not be suitable for the half-notes, which are played for most of the song. As always, there are two main options here: either compromise on a setting that will be neither perfect nor horrible on both note lengths, or simply automate the release so all the hits feel right.

These two tracks demonstrate the effect a release can have on the groove of rhythmical material. The effect might be easier to comprehend if one tries to conduct or move the head in relation to the beat.

Track 17.50: Drums Release Groove I
The long release on this track results in a gain reduction that does not track the groove of the drums. The relatively slow gain recovery interferes with the natural rhythm.

Track 17.51: Drums Release Groove II
A shorter release setting yields compression that is much more faithful to the rhythm.

Plugin: Digidesign Smack!
Drums: Toontrack EZdrummer

In the case of percussive performance, we want gain reduction to fully drop before the next note arrives.

Track 17.52: Snare II Uncompressed
In this uncompressed version, the velocity of each snare hit is identical.

Track 17.53: Snare II Release 1,000 ms
With 1 second of release, the gain never fully recovers between the various snare hits. Notice how the first hit is louder than all other hits, and how the last hit, which is the closest to the hit preceding it, is the quietest of them all.

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack EZdrummer

Peak vs. RMS

Compared to peak sensing, an RMS compressor responds more slowly to changes in the input signal levels. In a way, the RMS averaging is a form of attack and release that is applied before the threshold stage. Having a strong link to the way we perceive loudness, RMS is very useful for instruments with a less percussive nature, such as vocals. However, their slow response makes them somewhat less suitable for percussive instruments, which change in level very quickly. If *catch me if you can* is the name of the game, a compressor will do a better job in peak mode.

Peak can work better on a percussive performance.

Since the RMS function slows quick level changes, the whole compression happens in a more relaxed manner that makes the compressor less evident. Peak level detection leads to more obtrusive operation and often more aggressive results. If aggression or power is what we are after, peak mode can be the right choice. If transparency or smoothness is our aim, RMS comes in handy. That being said, some RMS compressors can be very aggressive, powerful, and punchy, while a compressor in peak mode can also do a decent job on vocals.

RMS for transparency; peak for a firmer control and more effect.

Peak-sensing allows the compressor to respond much quicker to transients. This might lead to extra loss in dynamics, for example, when compressing drums. Notice the loss of dynamics, attack, and punch in the peak version.

Track 17.54: Drums RMS

Track 17.55: Drums Peak
In the tracks above, no settings have changed apart from the peak/RMS parameter. This makes the comparison somewhat unfair since the peak version would natively present a harder, more evident compression. The next track involves the same settings as in Track 17.55, only with a longer attack; the result is very similar to the RMS version.

Track 17.56: Drums Peak Longer Attack

Plugin: Cubase Compressor
Drums: Toontrack EZdrummer

Side-chain control

The interaction between the threshold, ratio, attack, and release gives us extensive control over the compressed signal. The ability to alter the side-chain signal adds further control and a few more compression applications. The side-chain signal can be altered via two facilities: either the built-in side-chain filters or any processors we apply on the external side-chain. Applying processors of our choice on the external side-chain gives us ultimate control. Yet, in some cases, the built-in filters suffice.

The nature of the signal feeding the side-chain can be one of two: the original input, which we can equalize, or a different signal altogether. The latter is used for specific applications or as part of a unique effect. More on this soon. For now, let us see the benefits of modifying the original input.

One of the main characteristics of compressors is that they are more responsive to low frequencies. Put another way, low frequencies tend to excite compressors more than high frequencies. One reason for this is that due to their long period, low-frequency signals spend more time above the threshold. Another reason is the fact that low frequencies have to be louder in order to be perceived as loud as higher frequencies (as per the equal loudness curves). This characteristic of low frequencies is the very reason why compressors have a reputation for dulling sounds—as low frequencies are being compressed, the level of high frequencies is also reduced.

Low frequencies tend to trigger compression more than high frequencies, a fact that can yield dulling of compressed instruments.

To give a practical example, when overheads are compressed, the kick is most often the first drum to trigger the compression. So, whenever the kick hits, the cymbals are also reduced in level. In addition, if the threshold is set too low (or the ratio too high), we could defeat the kick. However, settings that are appropriate for the kick might not trigger any compression when the cymbals are played. Effectively, not all of the drum kit is being compressed uniformly. A typical consequence of this is crashes that drop in level in a steplike fashion every time the kick or the snare hits (this is often something we check when compressing overhead, and both the attack and release can affect this behavior). The same issue exists with any type of broadband sound—low and high frequencies are not compressed to the same extent. The way to solve this is to apply some low-frequency filtering on the side-chain. By applying, for example, an HPF, we make the compressor response more balanced to broadband signals such as overheads.

A high-pass filter on the side-chain can bring about more uniform compression of broadband material.

Track 17.57: Drums II Uncompressed
The source track for the following samples.

Track 17.58: Drums II No SC Filter
The compressor is set to the fastest attack and release in order to draw the most obvious effect. Both the kick and snare trigger heavy and fast changes in gain reduction, which yields fluctuation in the crashes' level.

Track 17.59: Drums II SC Filter
A 24 dB/oct HPF at 2 kHz is set to filter the side-chain. This essentially means that the kick does not trigger gain reduction anymore. Notice how the crashes' level is smoother as a result. As one would expect, filtering the low frequencies from the side-chain yields less gain reduction, so in this track the threshold was brought down and the ratio up, so the gain reduction meter closely hits the same 10 dB as in the previous track.

Track 17.60: Drums II SC Only
This is the filtered SC signal that triggered the compression in the previous track.

Plugin: McDSP Channel G
Drums: Toontrack EZdrummer

Then there is also an issue with a performance that varies in pitch. We can give the example of vocals. The low-mids of a vocal performance vary noticeably with relation to the pitch being sung. As the melody travels up and down, the lower notes trigger more compression.

The higher-mids of a vocal performance tend to provide more solid energy, which is less dependent on the sung pitch. Applying an HPF on the vocals' side-chain can yield more consistent compression. Perhaps the best demonstration of this problem has to do with bass guitars. The low E on a bass guitar (41 Hz) produces radically more low energy than the E two octaves above (165 Hz); therefore, the low E triggers more compression. Notwithstanding, the higher E will be more defined as its fundamental and harmonics start higher on the frequency spectrum. Having the low E less defined yet more compressed suggests how imbalanced the bass might sound. We can give the same example with the violin, trumpet, or virtually any instrument—the lower the note, the more the compression. Equalizing the side-chain can improve the compressor response to different notes. The actual filter being applied is often a high-pass filter (which can be of any slope), but low-shelving equalizers can also be suitable for this purpose. Sometimes even a parametric EQ will work.

Side-chain equalization, more specifically low-frequency filtering, can yield more uniform compression for different notes.

There is another way to look at the above—if all frequencies below 500 Hz are filtered from the side-chain, the same frequencies are likely to be louder as they trigger less compression. If we boost 4 kHz on the side-chain, 4 kHz will be attenuated more (this is de-essing, in essence). This leads to an interesting phenomenon—equalizing the side-chain affects the tonality of the compressed signal. Generally speaking, one should expect the opposite effect to the settings dialed on the side-chain EQ. But the applied tonal changes are far more sophisticated than simply being the opposite of the side-chain EQ—it is compression-dependent equalization, otherwise a form of **dynamic equalization**. Apart from having to consider how the side-chain equalization affects the tonality of the compressed signal, we can use it to fine-tune the frequency spectrum of instruments. It takes experience to master this technique, but it could be worth the effort.

Side-chain equalization alters the tonality of the compressed signal and enables dynamic equalization that is more sophisticated than the standard static equalization.

Applications

The following is a summary of compressor applications:

- making sounds bigger, denser, fatter;
- accentuating the inner details of sounds;
- adding warmth;
- containing levels;
- balancing levels;
- loudening;
- softening or emphasizing dynamics;
- reshaping dynamic envelopes;
- adding punch;
- accenting the natural attack;
- accenting transients;
- reviving transients;
- reconstructing lost dynamics;
- emphasizing or de-emphasizing decays;
- moving instruments forward or backward in the mix;
- making programmed music sound more natural;
- applying dynamic movement;
- ducking;
- de-essing; and
- applying dynamic equalization.

Along with the new applications listed here, some of the applications already discussed are covered in greater depth in the coming sections.

Accentuating the inner details of sounds

Compression invariably condenses the levels of signals. One reason we find the effect so appealing is that it tends to emphasize the inner details of sounds. The small nuances of any performance—the lip movements of the vocalist or the airflow from a trumpet—become clearer, which is largely satisfying. Like anything in mixing, over-doing can have a countereffect. It is like luminance: in the right amount, it can enhance an image and make its details clearer, but too much of it results in blinding white.

Balancing levels

Figure 17.33 shows the possible levels of three mix elements before and after com-pression. We can see that before compression, the level fluctuations of the different instruments are quite problematic: at points the snare is the loudest, at points the bass, at points the vocals; the snare level varies; and sometimes the bass is louder than the vocals, sometimes vice versa. It's a complete mess. No matter how much we try to
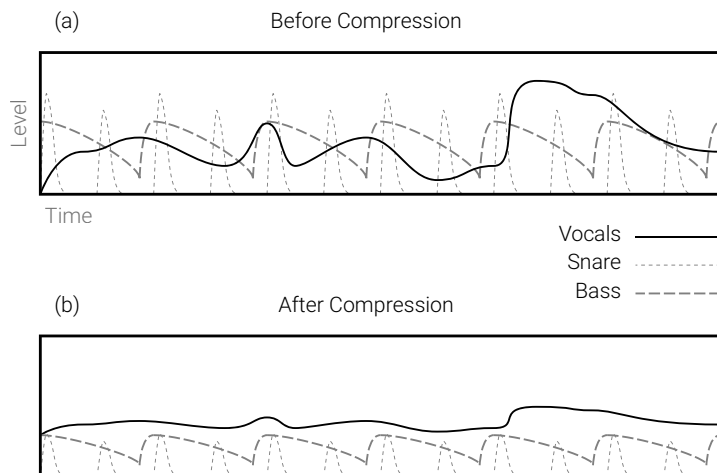
Figure 17.33 Relative levels before and after compression. (a) Before compression, the fluctuations in level make these three instruments impossible to balance. (b) After compression, the balancing task is easier.

balance such an arrangement using faders, the relative levels will always change over time and will do so quite radically. However, after we compress each instrument, we can balance them so their relative levels are consistent. Moreover, we can make them closer in level to one another, giving the impression of a louder, more solid mix.

Balancing of this sort usually happens very early in the mixing process, so as to make the tracks more manageable and instill some sense into their relative levels. Usually, we are not very fussy at this stage about perfecting the compression—rough settings would suffice to give reasonable balancing. As we progress from coarse to fine, we give more attention to the individual compression applied on each instrument.

Out of the three instruments shown in Figure 17.33, balancing the vocals is of prime importance as they exhibit the wildest level variations. Perhaps when working on the snare, balancing the level might be the second objective after adding punch. The level balancing approach of each instrument is likely to be different—we might want the vocal compression to be transparent while drawing more effect from the snare compression. If we have to generalize about sensible settings for level balancing, the following guidelines provide a possible starting point:

- **Threshold**—set at the bottom of the action range, so all level fluctuations are captured.
- **Ratio**—dependent on the extent of overshoots and to what degree we want to retain dynamics. A high ratio can squash an instrument, while a low ratio might not be sufficient to balance the levels. In Figure 17.33, chances are that the highest ratio will be applied on the vocals as they vary in level quite noticeably. Also, vocals have more settled dynamics, so softening these will be less noticeable.
- **Attack and release**—set to fast for less-percussive instruments, so more compression is applied and level changes are quickly tracked. At the same time, we must observe low-frequency distortion and how the natural attack and decay are affected. For percussive instruments, we have to set a longer attack and release to keep the dynamics.

- **Hold**—if available, can help reduce low-frequency distortion while letting us shorten the attack and release times.
- **Knee**—soft, as mostly when balancing levels we are not after the effect of compression.
- **Level detection**—depends on the instrument. Probably RMS for the vocals and peak for the snare.
- **Side-chain filter**—likely. Attenuating or filtering the side-chain lows of all three instruments can bring about more accurate results.
- **Makeup gain**—set so A/B comparison is fair (i.e., the perceived level will not change as we bypass the compressor).

One decision we often have to make is to what extent balancing should be done. If all the snare and kick hits are at the same level, they could sound like the product of a drum machine. Fortunately, the timbre of a drum changes in relation to the hit velocity, so some natural stamp is retained even if the individual hits are even in level. Yet, a more natural production could benefit from a relaxed balancing approach. In more powerful productions or more powerful sections of productions, many mixing engineers tend to even out the drum hits in a near mechanical fashion. We have to remember that when it comes to drums, there is a fine line between natural and sloppy. A good drummer will exhibit some consistency, some variations, and some consistency of variations. But if each hit is random in level, the performance comes across as amateur. As less sloppy means more even, perfectly even drum hits can give quite a powerful impression, albeit not necessarily a natural one.

> *Drum hits, even at velocity, can give the impression of a powerful performance, although this might not be suitable for natural genres.*

Track 17.61: Snare HH Unbalanced
The uncompressed source track, with loud snare hits followed by softer ones.

Track 17.62: Snare HH Balanced
Only the snare is compressed on this track. It involves a threshold at –30 dB, 4.4:1 ratio, and soft-knee. The attack, hold, and release are adjusted in order to try to maintain the timbre of the loud hits. The attack ended up at 0.7 ms, hold at 300 ms, and release at 290 ms. Also, a side-chain shelving filter has been applied to attenuate the lows, along with a parametric 4 dB boost around 3 kHz that controls the degree to which the louder hits are reduced in level. The gain on the parametric side-chain filter was used for the final balancing of the hits.

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack EZdrummer

Loudening

Loudening is achieved by increasing average signal level while maintaining its peak. During mixdown, this is done less in the mastering sense of maximizing the overall loudness of a stereo mix, and more with the objective of condensing the dynamics of

individual instruments, which often creates a more powerful impact. If a specific instrument fails to compete with other instruments in the mix, some compression can make it cut through. Loudening is also a means of balancing levels:

- **Knee**—soft-knee is obligatory here as the compression effect caused by a hard-knee will limit our ability to dial the extreme settings on other controls. The lower the knee span, the more evident the compression, but very high knee settings will limit our ability to push the levels higher. We generally try to have a knee span as wide as the action range of the input signal; 30 dB span can be a good starting point.
- **Ratio**—set to maximum. The soft-knee will provide gradual limiting, and the threshold control is used to determine the degree of true limiting.
- **Threshold**—when the threshold is set to half the knee span below 0 dB, the full knee effect is reached at 0 dB, and only signals above this level will be limited. Any additional lowering of the threshold will result in an identical extent of limiting range for loud input signals. For example, with a 40 dB soft-knee span and the threshold set to –20 dB, limiting will affect input signals at 0 dB and above (Figure 17.34a). Then bringing the threshold 10 dB down to –30 dB will yield limiting to the top 10 dB of the input signal (Figure 17.34b).
- **Makeup gain**—identical to the threshold, but inverted in magnitude. For example, for a threshold at –12 dB, the makeup gain would be 12 dB. Since the ratio is set to maximum, this matched setting brings the maximum output level to 0 dB.
- **Attack and release**—just like with balancing levels, short for less percussive and longer attack for more percussive. Low-frequency distortion must be observed.
- **Hold**—as with balancing levels, can reduce low-frequency distortion.
- **Level detection**—as we really need as little compression effect as possible, RMS is favored here.
- **Side-chain filter**—likely. Attenuating or filtering the side-chain lows can bring about more accurate results and reduce pumping.

(a)	(b)	(c)
Knee span: 40 dB	Knee span: 40 dB	Knee span: 40 dB
Ratio: ∞:1	Ratio: ∞:1	Ratio: 4:1
Threshold: 20 dB	Threshold: 30 dB	Threshold: 20 dB
Makeup gain: 20 dB	Makeup gain: 30 dB	Makeup gain: 15 dB
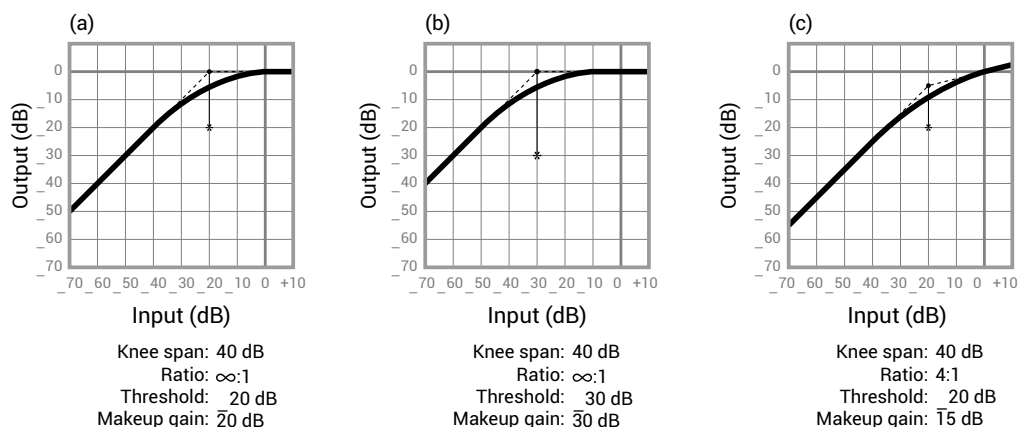
Figure 17.34 Three approaches to loudening. (a) The knee is reached at 0 dB. Any input signals above 0 dB will be limited. (b) A more aggressive compression is achieved by bringing the threshold down and the makeup gain up. The full knee is reached at –10 dB, so limiting starts for any input signals above –10 dB. (c) A softer effect can be achieved by lower ratio settings. This approach retains the dynamics of the input signal better than the other two.

This specific approach involves increasing compression on louder signals. Although it can result in extra power, it does not excel at retaining dynamics. We can tweak the different settings to achieve less-aggressive compression that will retain dynamics better. For example, if we can reduce the ratio to 6:1, the makeup gain will have to be brought down to compensate for the extra output level (Figure 17.34c).

Track 17.63: Bass Loudening I
This is a compressed version of Track 17.38. The threshold was set to –20 dB, the ratio to the maximum of 1,000:1, and the soft-knee span to 40 dB. The attack, hold, and release were all kept as short as possible—longer attack would mean higher signal peaks (which would limit the makeup gain) and longer release would mean less compression and less loudness. Notice the loss of bass attack compared to the source track. The makeup gain is set to +15.4 dB in order to match the peak reading of the source and this track.

Track 17.64: Bass Loudening II
Same settings as for the previous track, only with the threshold down to –30 dB, the ratio to 2.5:1, and the makeup gain to +5 dB. This track is slightly quieter than the previous one, but involves a bit more attack.

Plugin: Sonnox Oxford Dynamics

Reshaping dynamic envelopes

The level changes a compressor applies reshape the dynamic envelope of sounds. In the most simple terms, the threshold defines which part of the envelope is reshaped and the ratio defines to what extent it is reshaped. Attack and release give us far more control over envelope reshaping—we can alter the natural attack and decay in different ways.

The first application we will discuss is **adding punch**. Before showing how a compressor can achieve that, we should define punch. The definition of punch in dictionaries includes words such as thrust, blow, forcefulness, and power—all words that are not exactly associated with slow motion. They have to be abrupt to have an impact. Many of us know how restful sounds become as we open the release on a synthesizer. Staccato for frisky, legato for relaxed. One of the basic principles of music is that fast equals energy. It is hard to write a mellow love song at 160 bpm, just as it is hard to write a dance tune at 72 bpm. Consequently, as the tempo gets faster, the notes get shorter.

How does this relate to compression? We can shorten sounds using a compressor by attenuating their decay. It should be quite obvious by now how to do this—we set the attack to let the natural attack pass through and a long release in order to attenuate the natural decay. Figure 17.35 illustrates this on a snare. We can enhance the effect by shortening the attack, so some compression catches the natural attack, condensing it to give more impact.
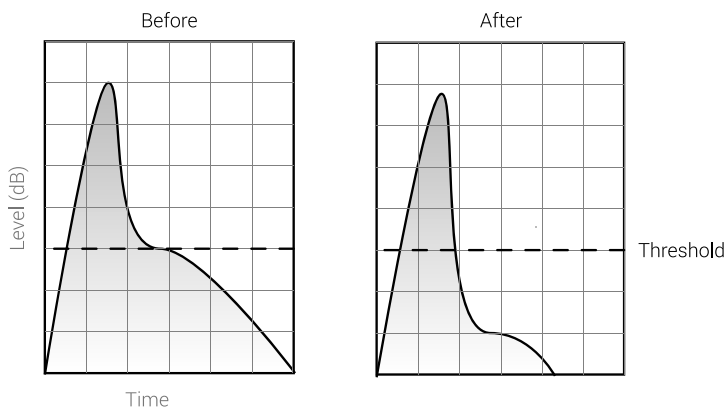
Figure 17.35 Adding punch to a snare using a compressor. The attack is set long enough to retain the natural attack, while the release is set long so gain reduction brings down the natural decay.
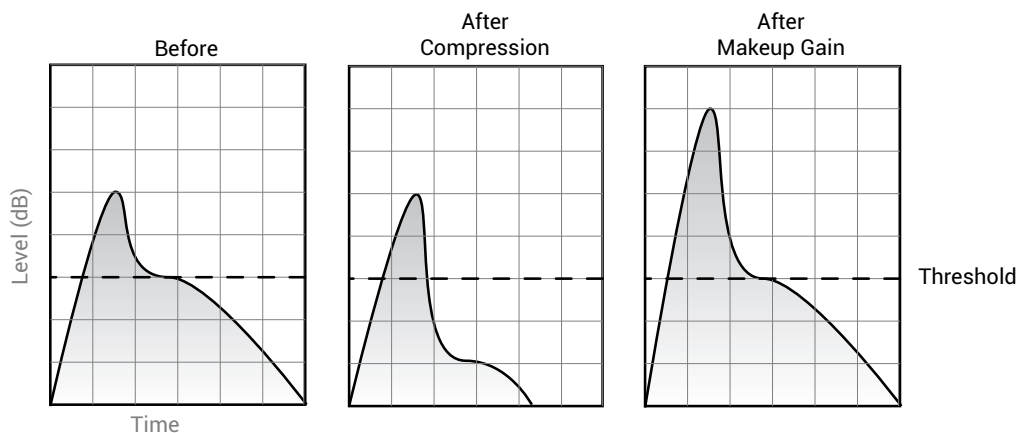
Figure 17.36 Accenting the attack of a snare using a compressor. The attack is set long enough to retain the natural attack, while the long release brings down the decay. After applying makeup gain, the attack is made louder, but the decay returns to its original level.

We can use exactly the same principle to **accent the natural attack**, **accent transients**, **revive transients**, or **reconstruct lost dynamics** (due to over-compression, for example). All we need to do is apply some makeup gain after bringing down the decay. Figure 17.36 illustrates this. We can see that after the addition of makeup gain, the attack is emphasized compared with the input signal before compression.

There is more in the dynamics reshaping examples above than first meets the eye. Compression is a sophisticated way of **bringing instruments forward or backward** in the mix. Perhaps we want to bring the kick forward, but raising the fader detaches it from the mix and makes it too dominant. On the same basis, we might want to pull the kick back a bit, but even gentle lowering of the fader results in a faint kick. One of the problems with using faders is that they affect the full frequency spectrum of sounds. Consequently, this affects the frequency balance of the mix and the masking interaction between instruments. If we want the kick forward in the mix, why not just bring up

its attack? If we want it farther away, why not just soften its attack or perhaps its decay? Only affecting a part of the dynamic envelope means we do not have to consider issues such as masking as much. Earlier, we saw how equalizers can provide a surgical alternative to faders by only affecting limited frequency ranges in sounds. Compressors (and other dynamic range processors) can achieve something similar, except they affect limited time (or length) of sounds. In essence, compressors can be used for sophisticated level adjustments as a kind of surgical faders.

Compression can be used to bring instruments forward and backward in the mix.

Track 17.65: Reshaping Source
The source drums used in the following samples. The snare in this track is not compressed.

Track 17.66: Reshaping More Attack
For the purpose of demonstration, the attack of the snare is accented to an exaggerated degree. First, the ratio was set to the maximum of 100:1, both time constants were set to minimum, and the threshold was brought down to –40 dB. Then the attack was lengthened, and the longer it became, the more of the natural snare attack passed through; it ended up at the maximum value of 300 ms. (This extremely long attack was required due to the extreme gain reduction of around 15 dB. There are less than 200 ms between the first and second hit, but remember that the attack time determines the rate—not the time—it takes gain reduction to rise.) The release was then lengthened to 1 second so as to track the natural snare decay. The gain reduction was set to +11 dB in order to bring the snare decay back to its original level. Notice how the snare has more punch now, and how it has moved forward in the mix. Also, notice that, as opposed to the previous track, the attack of each hit is not consistent in level. This is due to the release not recovering fast enough between the hits. The sooner each consecutive hit arrives, the less attack it has. For example, the first hit has the loudest attack, and the one straight after it has the softest. There is little we can do about this since shorter release means that the decay of the snare would alter; either bringing the threshold or the ratio down in order to reduce the amount of gain reduction (and thus the quicker the gain would recover) would mean less natural attack compared to decay. We will soon see how this can be resolved using another compressor.

Track 17.67: Reshaping Less Attack
The settings involved in reducing the attack of the snare are completely different from the ones used in the previous track. This time, we only want the compressor to reduce the natural attack without affecting the decay. The threshold is set to –13 dB, a level above which the snare only overshoots by 5 dB. The ratio is set to 3.2:1. Both the attack and release are set to minimum –14 μs and 5 ms, respectively. No makeup gain is applied. The result of this compression is a snare with less punch, positioned farther back in the depth field.

Track 17.68: Reshaping Reviving Attack
For the sake of demonstrating how a compressor can revive lost transients, the attack-lacking snare from the previous track is fed through another compressor with the same settings as in Track 17.66. It is hard to miss how it has regained some healthy attack.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3
Drum: Toontrack EZdrummer
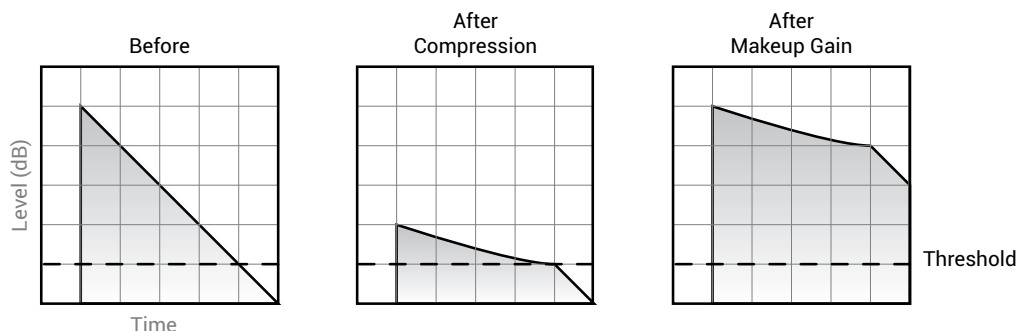
Figure 17.37 Emphasizing the decay of a reverb. The low threshold, high ratio, and fast attack reduce the reverb quickly and drastically. As the reverb decays, the gain reduction recovers quickly due to quick release settings. After applying makeup gain, the reverb decay becomes evidently louder.

One case we have not touched upon yet is how compression can be used to **emphasize decay**. Although compression brings levels down, we have already seen how we can accent the natural attack by attenuating the decay then raising the makeup gain. To emphasize decay, we attenuate the attack but leave the decay as is, then the makeup gain brings the whole signal up. This can be extremely useful when we want reverbs to be more present. The compressor settings in this case involve low threshold, high ratio, fast attack, and fast release so the gain reduction recovers as quickly as possible. Figure 17.37 illustrates this. The very same settings can be used to **lengthen the sustain** of instruments, notably guitars. It should be noted, however, that compressing natural reverbs is known to impair the spatial information within them—a faithful sense of space is dependent on the reverb's natural dynamics (notably its decay). If reverbs are to appear natural, extra care has to be taken while compressing them. Conversely, when reverbs are used in a creative context, compression can easily improve the effect.

While percussive instruments are a good source for demonstrating the attack and release controls, we should not forget that microdynamics control can also be useful for less percussive instruments. Taking an acoustic guitar as an example, sometimes its sound is too jumpy due to, say, the use of a small diaphragm condenser microphone. This is not in the sense that it fluctuates in levels; it might simply sound "liver than life." We can use a compressor to smooth spiky microdynamics and make the overall sound more rounded.

De-essing

One way to treat sibilance is to attenuate the voice whenever sibilance occurs. This is done by feeding into the side-chain a signal that rises significantly in level whenever sibilance occurs. To achieve this, we boost the sibilance frequency on the side-chain signal, while often also filtering frequencies below it. Whenever the sibilance peaks, the side-chain signal overshoots the threshold and the compressor pulls down the overall voice. One issue with this manual approach is that it attenuates not only the sibilance, but the rest of the voice as well, including its body and higher harmonics. Another issue is that the sibilance on low-level passages might not be loud enough to overshoot the threshold, so only the sibilance on loud passages is treated. A well-designed dedicated de-esser takes into account these two issues, so it is most likely to do a better job.

Making programmed music sound more natural

The dynamic complexity of real performance is beyond comparison to what any computer can simulate. To give an example, synthesizers have attack and release controls, but none of them affect a chord already being played. On a real piano, the sound of every chord alters, and is being altered by, the chord succeeding it. Another example would be a crash hit—a resting crash being hit would sound different from a hit on a crash already swinging. Snares and many other instruments behave very much the same way. If a programmed musical event is a note, chord, or drum hit, the dynamics of a programmed sequence are applied on each event at a time, but there is no dynamic link between the different events. Compression works on all events as a whole; thus, it can be used to link them in a real-life fashion.

A classic example would be the machine-gun-like sound of programmed snare rolls. Such mechanical arrangements cause each hit to stand out as an individual hit. We can blur the gaps between the hits and meld them together using a compressor. Generally, we want fast attack to make each hit stand out less, and fast release so the decay is retained.

Applying dynamic movement

We have discussed how useful compressors can be for balancing the varying levels of a performance. Unless used to reshape dynamic envelopes, compressors are often associated with the restriction of dynamics. But it is not a secret that a mix with static dynamics can be quite boring. Compressors can be used to get in some action.

To give an extreme example, we can consider an organ playing harmony. The keyboard player changes a chord each bar and all the chords are at reasonably the same level. We can feed this organ into a compressor and feed the kick into the side-chain, and then set low threshold, moderate ratio, fast attack, and a release that will cause full gain recovery by the time the next kick hits. Every time the kick hits, the compressor will attenuate the organ, which will then rise in level until the next kick arrives. It is automatic automation. By attenuating the organ with every kick hit, we also clear some frequency space for the kick. We can also do something similar in order to create more sophisticated drum loops. We can go even further by adding synced delay to an external side-chain signal, so the dynamic movement is triggered by one event, then locks to the tempo.

A compressor as a ducker

A compressor used in the way described above—to attenuate a signal in relation to another signal—behaves like a ducker. A compressor is not really a ducker as the amount of attenuation a compressor applies is dependent on the overshoot amount, whereas a ducker applies a fixed amount. With a compressor, the attenuated signal fluctuates in level in relation to the side-chain signal, which is not the case with duckers. Figure 20.5 illustrates these differences. Percussive instruments are generally less of an issue, since they do not tend to fluctuate as much above the threshold. Also, we can always lengthen the release to slow down gain-reduction changes, but this might not yield musical results once the side-chain signal drops below the threshold. Although for steady ducking it is best to use a ducker, compressors can be, and are, employed as well.

Figure 17.38 Making a compressor behave like a ducker. The vocal track is sent to a bus called Gate SC. The moderate track has a signal generator on its first insert slot (generating pink noise), followed by a gate that has its side-chain input set to the Gate SC bus. This moderate track is not mixed, but routed to a bus called Gated Noise instead. On the distorted guitar track, a compressor is loaded, with its side-chain set to the Gated Noise bus.

A full exploration of the fascinating ducking applications is presented in Chapter 21. For now, let's look at a small trick that can turn any compressor into a real ducker. We are trying to prevent the gain reduction from fluctuating in relation to the input signal. We could use long hold or release times, but we need these to control the gain recovery once the side-chain signal drops below the threshold.

We have to employ a moderate track. Say we want the vocal to duck the guitars. The moderate track is fed with pink noise, which is gated in relation to the vocal track. We set the gate to open every time the vocal crosses a specific threshold. The moderate track is not mixed to the main mix, but instead routed to the side-chain of the compressor on the guitar track. When the vocal overshoots, the gate opens and the pink noise will trigger the compression of the guitar. The pink noise can be regarded as steady in level, so the guitars will always be attenuated by the same amount. Figure 17.38 shows a snapshot of this setup.

Tricks

Parallel compression

Parallel compression is one of the oldest and more common tricks in mixing. Parallel compression had already been implemented as part of the Dolby A noise reduction system, which was introduced in 1965. In 1977, Mike Bevelle published an article in *Studio Sound* magazine, which gave this technique much of its publicity. Bob Katz coined the term *parallel compression*, with other names being *side-chain compression* (which is confusing since it is not the compressor side-chain we are compressing) and

the *NY compression trick* (NY engineers were notorious for using this technique). It is usually applied on drums, bass, and vocals, although it can be applied on virtually any instrument.

The compression we have discussed so far in this chapter is downward compression—loud signals are brought down in level. One fundamental problem with this practice is that our ears are more sensitive to the bringing down of loud signals than they are to soft signals being brought up. Also, the compression artifacts tend to be more noticeable as they are triggered by loud level signals.

The idea of parallel compression is quite simple—instead of bringing high levels down, we bring the quiet levels up. The ear finds this all the more natural, and the artifacts triggered by low-level signals make the compression effect less evident. In addition, parallel compression retains dynamics much better than downward compression as transients are not brought down—if anything, their bottom is beefed up. Being more transparent, parallel compression lets us drive the compressor even harder when we are after a stronger effect. Also, we can adjust the degree of impact using a single fader rather than a multitude of compressor controls.

> *Parallel compression lets us make sounds even bigger while retaining their dynamics.*

Parallel compression involves a simple setup: take a copy of a signal, compress it, and blend the compressed version with the uncompressed original (Figure 17.39). Generally speaking, the original remains consistent in level and the compressed version is layered below it at a level of our choice.

There are various ways to achieve that within an audio sequencer. If it is a single track in question, we can simply duplicate the track and compress the duplicate. However, if we change something on the original track (e.g., the EQ), we also have to change the duplicate. A more popular approach involves sending the track in question—or a group of tracks—to a bus; the bus is fed into two auxes, and a compressor is only loaded onto one aux (Figure 17.40). We then alter the level of the compressed aux track. We must make sure that there is no delay between the original and the compressed versions or comb filtering will occur—auto delay compensation ensures this will not happen.
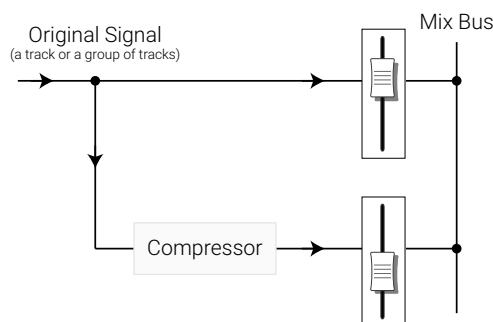
Figure 17.39 Parallel compression is a blend between a compressed and an uncompressed version of the same source. The level of the compressed version is what we alter for more or less effect.
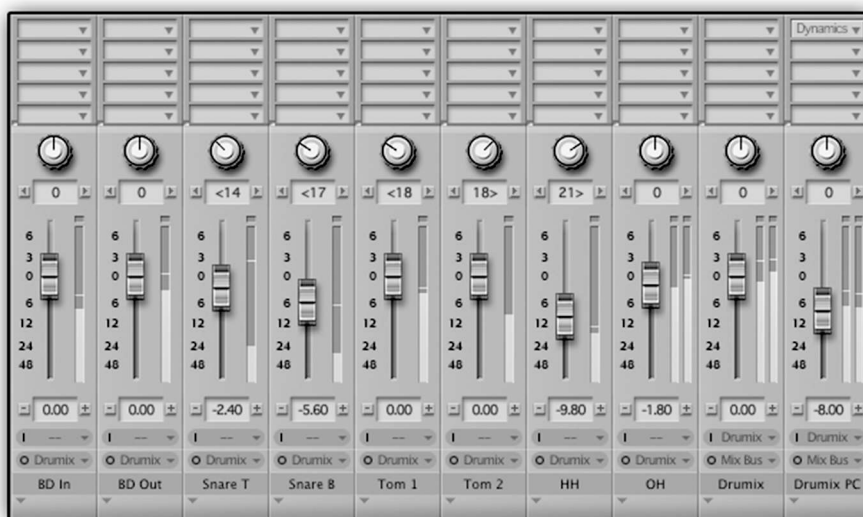
Figure 17.40 Parallel compression in Digital Performer. The output of various drum tracks is routed to a bus called "drumix." The bus feeds two auxiliaries—drumix and drumix PC. On the latter, a compressor is loaded, and it is blended lower in level with the drumix auxiliary.

While additional routing is required when parallel compression is done on an analog desk, at least one aux track in an audio sequencer could be spared if compressor plugins were given the ability to mix the compressed and the uncompressed input signals (something very similar to the standard wet/dry control on a reverb emulator). While this requires very little effort to implement, only a few compressors provide such a control. The PSP *Vintage Warmer 2* in Figure 17.41 is one of them.

Parallel compression is a form of upward compression—instead of making louder sounds quieter (downward compression), we make quiet sounds louder. But parallel compression is different from upward compression in that low-level signals are boosted in a linear fashion (1:1 ratio), the most drastic ratio happens right above the knee, and from there up the ratio curve slowly returns to 1:1. Figure 17.42 illustrates this. We can consider parallel compression as having the most effect on medium-level signals as they cross the threshold. In practice, however, the threshold is set quite low below the action range, so it is eventually the quiet signals that get most of the effect. In essence, the quieter the signal is, the more it is brought up in level; high-level signals are hardly boosted. This can be very useful for level balancing.

We can see from Figure 17.42a that the ratio transition around the knee can be quite drastic. We can achieve a softer transition in two ways: either bring down the ratio or bring down the level of the compressed version. As we bring down the compressed version in level, the resultant sum will morph toward the unity gain curve of the uncompressed version. With the compressed version all the way down, the resultant sum is identical to the uncompressed version.

The attack time of the compressor is set to be as fast as possible, so sharp transients are not overemphasized (unless this is what we want). If transparency is what we are

Figure 17.41 The PSP *Vintage Warmer 2*. The Mix control (to the right of the Brick Wall switch) lets us blend the unprocessed input with the compressed version, making parallel compression a very easy affair.
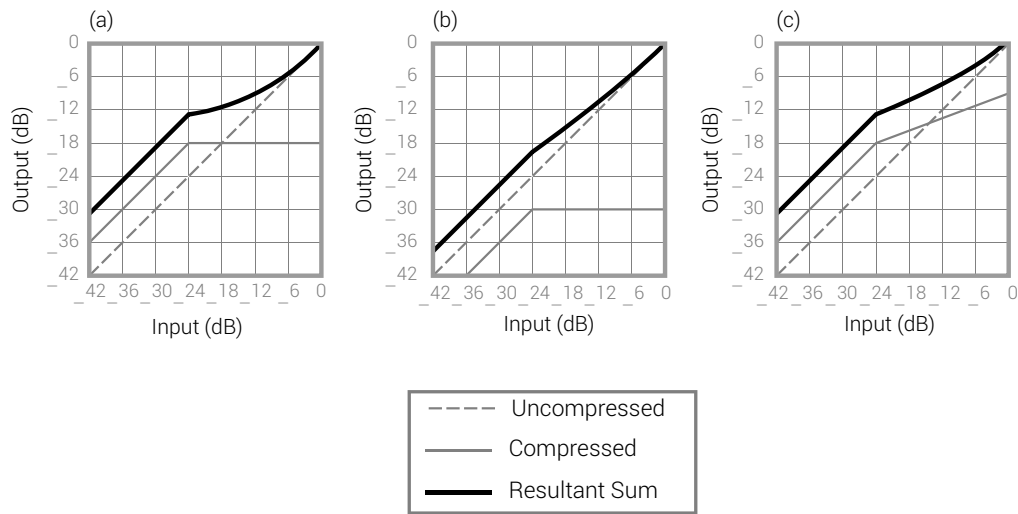
Figure 17.42 Parallel compression and resultant ratio curves. (a) The compressed version, with a 1,000:1 ratio, is 6 dB louder than the uncompressed version in order to demonstrate the shape of the resultant ratio curve. We can see that, from the 1:1 ratio below the threshold point, the ratio slowly returns to 1:1 as the signal gets louder. (b) The same settings as in (a), only with a compressed version that is 6 dB below the uncompressed version. We can see a smoother ratio curve above the knee. (c) The same settings as in (a), only with 2:1 ratio. Again, we can see a smoother ratio curve above the knee.

after, the release time is set to minimize pumping and low-frequency distortion. However, a powerful effect can be achieved if the compressed version is pumping and distorting, and, again, we control how much of this effect blends in using the level of the compressed version.

The parallel compression technique is almost matchless when applied to a group of instruments that has already been dynamically treated, such as a drum mix. Applying a standard compression on a drum mix can easily impair the dynamics of each source track. For example, while setting the drum mix compression, we have to find the right attack and release times that will retain the dynamic character of each drum. Parallel compression makes the whole affair of adding power less restricting.

Parallel compression is also one of the rare techniques where a compressor might be used in an aux send arrangement rather than the standard inserts. Taking a drum mix, for example, we can create an alternative submix using sends and feed the send-bus to the parallel compressor. Then, if we want a bit more weight on the kick, the kick send-level is brought up.

Track 17.69: PC Drums Uncompressed Layer
The uncompressed version of the drums.

Track 17.70: PC Drums Compressed Layer
The compressed version of the drums in the previous track. The compressor is set for heavy effect with the fastest time constants.

The following tracks are simply a blend of the two previous tracks. The uncompressed layer is always mixed at the same level, with the compressed layer always a few dB below it (denoted in each track name). Notice how much better the attack and dynamics are retained compared to standard compression:

Track 17.71: PC Drums 21 dB down

Track 17.72: PC Drums 15 dB down

Track 17.73: PC Drums 9 dB down

Track 17.74: PC Drums 3 dB down

The same set of tracks, but this time involving a bass:

Track 17.75: PC Bass Uncompressed Layer

Track 17.76: PC Bass Compressed Layer
The intentional distortion on this track is the outcome of fast time constants.

Track 17.77: PC Bass 21 dB down

Track 17.78: PC Bass 15 dB down

Track 17.79: PC Bass 9 dB down

Track 17.80: PC Bass 3 dB down

Plugin: Universal Audio Fairchild
Drums: Toontrack EZdrummer

Serial compression

Serial compression, or *multi-stage compression* as it is sometimes called, is also an old technique. It involves connecting two (or more) compressors in series. In an analog studio, this calls for extra routing and the loss of either a channel strip or an external unit. There is no such overhead with audio sequencers—setting up serial compression simply involves loading two plugins in consecutive insert slots.

The idea behind serial compression is this: If we have two tasks to achieve, why not distribute each task to a different compressor and have the settings on each compressor optimal for its task? For example, if we want to add punch to a snare, but not all the snare hits are at the same level, we can use the first compressor to balance the level of the performance and the second to add punch. This mix-and-match approach is also beneficial since some compressors are better for specific tasks—one compressor might be better at balancing, another at adding punch.

Track 17.81: Reshaping Balancing Attack
You probably remember the problem in Track 17.66, where the attack of the different hits varied in relation to hit spacing. To solve this, another compressor is inserted following the compressor that accented the attack. The balancing compressor has its threshold set at –18.3 dB, a ratio of 2.8:1, soft-knee, 15 ms attack, 5 ms release, and 3.2 dB of makeup gain.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3
Drums: Toontrack Ezdrummer

More on compressors

Compressors, depth, and focus

We have seen in the previous sections how manipulation of the dynamic envelope of instruments can send them backwards or forward in the depth field. Part of this phenomenon has to do with an increase or decrease in the overall level of the treated instrument. There is another important aspect of compression that links it to depth and focus. It involves recordings that include ambiance or reverb.

The relationship between reverbs, depth, and focus will be expounded in Chapter 24 (on reverbs). For now, it is sufficient to say that the more reverb there is compared with the dry sound, the farther away and the less focused sounds will appear. Compression can either increase or decrease the level of a reverb already printed on a recording. It is easy to demonstrate this with a drum recording that contains ambiance, such as overheads or a drum mix. If the compressor is set for short attack and release, the reverb will become louder, making the drums appear farther back and their image more smeared. If the compressor is set with an attack long enough to let the natural attack through, and longer release so some gain reduction is still applied on the ambiance in between hits, the drums will appear closer and their image will sharpen.

Experienced or attentive readers have probably noticed the effect of compression on depth and focus in virtually all previous samples that included a drum kit. For example, in Track 17.2, the drums move farther back as the threshold drops. Here are three more tracks demonstrating this.

Track 17.82: Drums III Uncompressed
The source track will be compressed in the following two tracks.

Track 17.83: Drums III Backwards
Short time constants and gain reduction peaking at around 7 dB are enough to noticeably increase the ratio of the ambiance level. This sends the drums backward in the depth field and makes their image less focused.

Track 17.84: Drums III Forwards
The long attack and moderate release in this track are set to let the kick and snare pass through but maintain some gain reduction after each hit, thus reducing the level of the ambiance. Both drums appear more frontal and focused.

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack Ezdrummer

Setting up a compressor

There are various ways to tackle the initial settings of a compressor. The method mentioned here provides a step-by-step approach that can be helpful in many situations, notably when compressing vocals or drums. People new to compression are the most likely to find this technique useful. It can also be used as an *aural exercise*—it makes the function of each control appear rather obvious. This method is based on three main stages—the extent, degree, and timing—which in relation to controls are threshold, ratio, and attack and release.

- **Initial settings**—the idea is to have the most obvious compression once the threshold is brought down in the next step. All controls but the threshold are set to produce the hardest possible compression.
- Threshold all the way up (so no compression takes place initially).
- Ratio all the way up.
- Attack and release as short as possible.
- If there is any option regarding peak/RMS or soft-/hard-knee, these are better set at this point based on the nature of the instrument. For the purpose of the aural exercise, peak and hard-knee are preferred.
- **Determine extent—threshold** goes down. Due to all the other settings, as soon as signals overshoot, the compression becomes very obvious and on many compressors the fast attack and release will cause audible distortion. We can easily identify what portion of the signal is being compressed and adjust the threshold accordingly. It can be beneficial to set the threshold slightly below what seems to be the appropriate setting, as later it can lead to more uniform compression.
- **Determine degree—ratio** goes down from maximum. Once we know the extent of compression, we can determine its degree.

- **Determine attack**—the attack is made longer from its initial short settings.
- **Determine release**—the release is also made longer.

There is a sense to the order in which controls are set with this method. Changes to the threshold and ratio will cause changes to the amount of gain reduction, which in turn affect the overall attack and release times (the more gain reduction, the slower the attack and release). So the attack and release are set after the threshold and ratio. The attack is set before the release, since long release settings might interfere with succeeding attack phases. While this method can get very close to what we need, usually fine-tuning each control will follow.

Here is a demonstration of the steps involved in this method, using the uncompressed vocal in Track 17.17. The compressor is initialized as described above, and to make the effect more obvious the knee is set to hard.

Track 17.85: Vocal Steps Threshold
The threshold was brought down, ending at –22 dB. This level is slightly lower than what might seem to be needed, to allow longer attack later on.

Track 17.86: Vocal Steps Ratio
The ratio is brought down from 1,000:1 to 3:1, at which the vocal is reasonably balanced.

Track 17.87: Vocal Steps Attack
One problem with the previous track is that the function of the attack was too obvious, so it has been lengthened to 1.85 ms. The release is left untouched at the minimum value of 5 ms.

Plugin: Sonnox Oxford Dynamics

Multiband compression

The compressor we have discussed so far is known as a broadband compressor—it attenuates the whole input signal, with its full frequency spectrum. We have already seen some issues this can bring about—it can cause the dulling of sounds due to high-frequency attenuation when low frequencies trigger the compression. Also, when we compress a bass guitar, low notes will result in more gain reduction, which also affects the mid- and high frequencies that give the bass guitar much of its definition. When we de-ess vocals, we attenuate not only the sibilance, but also the body and the higher harmonics of the voice.

A multiband compressor, or a *split-band compressor*, splits the input signal into different bands, and then lets us compress each band individually (Figure 17.43). For example, for de-essing, we can define a band between 3 and 8 kHz, and only sibilant peaks at this frequency range will be compressed, without affecting the rest of the frequency spectrum. We can employ the same principle when trying to de-emphasize the resonance of an acoustic guitar.

The manufacturing of an analog multiband compressor is more costly as it involves more than one compressor and a quality crossover. While being a standard tool in mastering studios, multiband compressors were never high up in the mixing priority
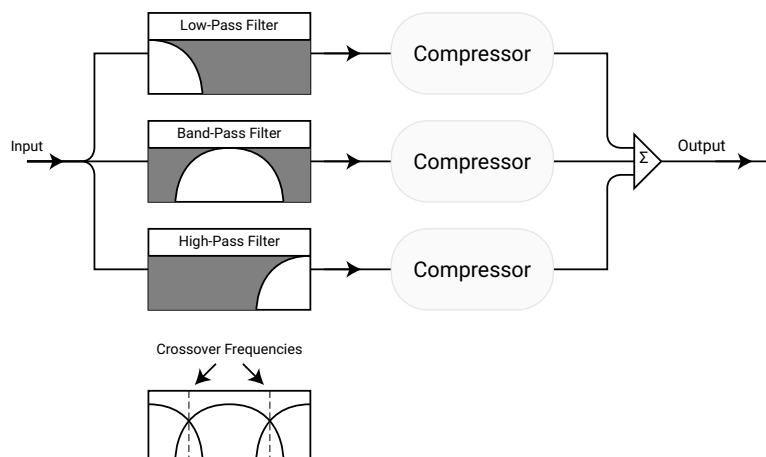
Figure 17.43 A block diagram of a multiband compressor. In this illustration, a three-band compressor is shown. The input signal is split into three, and filters are used to retain the respective frequencies of each band. Then each band is treated with its own compressor, and the three bands are summed together to constitute the output signal. At the bottom, we can see the combined effect of the filters and the crossover frequencies between the different bands.

list. However, software plugins cost nothing to manufacture, so these have made multiband compressors more widespread today (Figure 17.44).

Most multiband compressors have a predefined number of bands and they let us set the crossover frequencies between the different bands. A dedicated compressor per band should offer most related controls. Some controls might still be global and affect all bands.
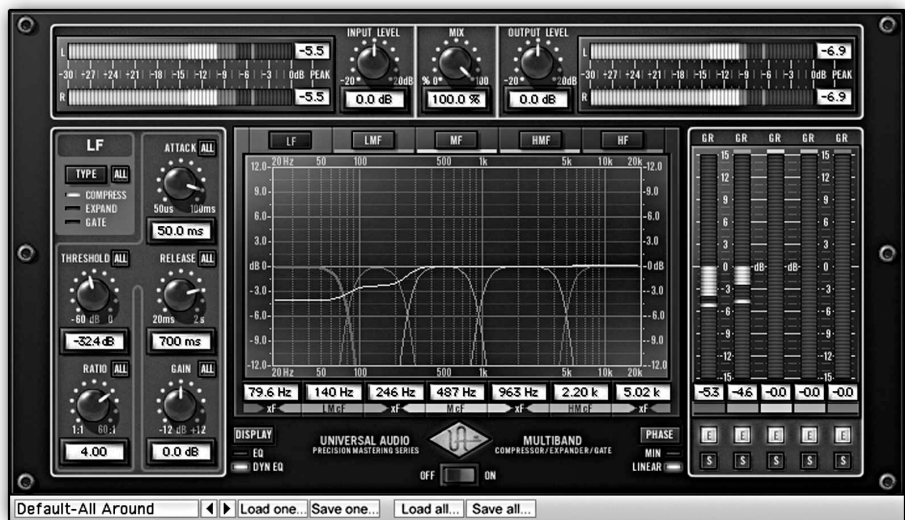
Figure 17.44 The Universal Audio *Precision Multiband*. This plugin offers a compressor, expander, or a gate for each of the five bands. On the left are the controls for the compressor of the selected LF. The plot in the middle shows the frequency division between the five bands. The meters on the right show the gain reduction per band.

Multiband compressors provide great advantages when the signal in question is of a broadband nature, such as a stereo mix. Compressing each band individually provides a greater degree of control over the signal. This is why multiband compressors are so popular in mastering. Nevertheless, every instrument being mixed can be divided into a few separate frequency bands that each plays a different role, and we can benefit from compressing each band individually rather than the instrument as a whole; for example, the lows of a kick and the attack on its high-mids. The softer the drummer hits the kick, the less are the lows compared with the click of the beater. We can compress the low band while compensating with some makeup gain specifically on that band. This will make the lows-to-attack ratio more consistent with different hit velocities. We can apply the same principle on a bass guitar in order to achieve a more solid frequency balance for different notes. We can also reverse this idea in order to make a pro-grammed kick more natural. Unless velocity layers are used, kick samples have the same ratio between lows and high-mids. We can set a low threshold and a low ratio on the high-mid band, so the higher the hit velocity is, the more the high-mid attack will be compressed. This will alter the ratio between the lows and attack as in the real world. Last but not least, multiband compression is an excellent way to rectify the proximity effect on vocal tracks, with the low-end boost resulting from the vocalist moving closer to the microphone. Compressing the low bands rectifies these variations.

Another advantage of using multiband compressors is that, while only affecting a specific frequency band, they tend to interfere less with the overall tonality of the processed signal. In addition, the compression on each band can be pushed harder since the compression artifacts are introduced on a limited range of frequencies. For example, fast release on the lower band will cause distortion, but we can dial a faster release on the higher bands, which will make these bands louder. By making each band louder separately, we can make signals even louder than we could with a broadband compressor. It is a divide-and-conquer approach applied to the frequency domain. Although they involve more effort and time to set (and a little more auditory skill), multiband compressors give us more control over the dynamic treatment of sound.

When a multiband compressor is not available, we can set up a manual arrangement using an array of filters and broadband compressors. Say we need two bands, we can send a track to a bus and bring the bus back into two different tracks, put a low-pass filter on one and a high-pass filter on the other, and then set both to the same cut-off frequency. Next, we insert a compressor on each track, and send both at the same level to the mix bus. If multiband compressors were as simple as explained so far, this manual setup would work flawlessly. However, multiband compressors involve some additional science in the way they sum the different bands back together. For example, they might align the phase of the different bands. Manual multiband compression does work, but a true multiband compressor should sound better.

Track 17.88: Bass Multiband
This track demonstrates how a multiband compressor can be utilized to reduce squeaks. In the original bass recording (Track 17.38), there are two noticeable nail squeaks on the A/G# notes, and then a few more toward the end. Squeaks as such might be masked when the guitar is mixed with other instruments, but they can become annoyingly noticeable if the

bass is equalized, distorted, or processed using an enhancer. The compressor is set with a single active band between 444 Hz and 3.3 kHz. Heavy gain reduction occurs on this band only when the squeaks happen. Compare this track to the original to see the difference.

Track 17.89: Bass Multiband SC
Most multiband compressors let us solo each band. This is the soloed version of the active band in the previous track, with its compressor disabled. Essentially, this signal feeds the side-chain of the compressor on the band used.

Plugin: Universal Audio Precision Multiband

Before or after the EQ?

What seems like an eternal mixing question is: Should a compressor come before or after an equalizer? Had there been a determined answer, this question would not be asked. The answer is: it depends. Let's see on what.

We have already established that low frequencies excite compressors more than high frequencies. There are additional situations where we clearly do not want the compressor to respond to a specific frequency range of the input signal, in which case an equalizer has to be placed before the compressor. For example, we might want to filter the lows of vocals or overheads before compressing them. The problem with a pre-compressor equalization is that any changes to the equalizer affect the characteristics of the compression. So, once the compressor is set, touching the equalizer might require a revisit to the compression settings.

So the wisdom is this: for *compression-related equalization*, we should put an equalizer before the compressor. Any other equalization should be done after the compressor. This includes equalization that alters the tonality of an instrument or one that compensates for the tonal alterations caused by the compressor.

*All but compression-related equalization should be done
after the compressor.*

Many compressors nowadays let us equalize the side-chain signal. When such a facility exists, equalizers are most likely to be placed after the compressor. There is a huge benefit of using side-chain equalization rather than pre-compressor equalizers. With pre-compressor equalizers, the tonality of the signal is affected by two equalizers—the one before and the one after the compressor. However, the pre-compressor equalizer is said to be locked for the reasons explained above. None of this happens with side-chain equalization—the tonality of the signal is solely controlled by a single equalizer placed after the compressor. Perhaps the only shortfall of using side-chain equalizers is that they often provide fewer bands, or less control, than a standard equalizer. These might not be sufficient for our needs, so pre-compressor equalizers might still be employed.

*Providing their facilities are adapted for the task, side-chain equalizers
provide great benefit over pre-compressor equalizers.*

Having said that, it pays to remember that post-compressor equalizers can damage the dynamic control achieved by the compressor. For instance, we might have balanced vocals out of the compressor, but then boosting around 1,760 Hz would make all the A notes jump out (1,760 Hz is A, and thus a harmonic of all A notes below it). There is no magic solution to such a problem. Sometimes the choice of a pre- or post-compressor equalizer is a matter of experimentation.

Compressors on the mix bus

Engaging a compressor on the mix bus while mixing, also known as *bus compression*, is another source of a long debate in the mixing community. Some mixing engineers use bus compression passionately from very early stages of the mix. Often they have at their disposal a respectable compressor, such as the SSL bus (or quad) compressor, and they drive it with rather subtle settings. The bus compressor might be used during the majority of the mix, but it is often switched out before the mix is printed—mastering engineers seem rather united against mixes that have already been stereo-compressed. Conversely, some mixing engineers would rather be mixing on the edge of an airport runway than mix with a compressor on the mix bus—they are that much against it. Both types include some of the world's most creditable mixing engineers.

Perhaps the least practical reason for using bus compression is that it simply makes the mix more exciting, and excitement is an important part of mixing. Bus compression also lets us evaluate how the mix might translate to radio or in clubs—both involve heavy compression and/or limiting in their signal chain. Mostly in these situations, we care how the low-end affects the rest of the mix. Is the kick too boomy? Is it thin?

There are a few risks in using bus compression. The greatest of them is that it limits our ability to judge relative levels. If the bus compressor is set to anything but subtle, it should be no wonder if, when it is bypassed, the mix relative levels turn chaotic. Also, by applying compression, signals that overshoot the threshold are reduced by a specific ratio. If a fader is brought up by 4 dB, its instrument might only rise by 1 dB due to a 4:1 ratio on the bus compressor. In the way compressors work, pushing the kick up might result in more compression of low frequencies, which can actually make the kick quieter (there is no individual makeup gain when the stereo mix is compressed). So it is not only the precision of the faders that we can lose with bus compression—level changes might be opposite to fader movements. Another thing to ask is how exactly one should go about servicing the bus compression. Any change in the mix can also affect the compression. So should we revisit the bus compressor after every mixing move? It should be clear that only respectable compressors and subtle settings would be fit for the task.

Perhaps above all comes the idea that only trained ears can recognize the subtleties of bus compression. For the novice, bus compression can be like a translucent layer that only blurs what is behind it—the actual mix.

Unless you have a good compressor, subtle settings, and trained ears, bus compression can be self-defeating.

Dynamic convolution compressors

The concept of convolution, in simple terms at least, involves capturing the characteristics of a specific audio system so they can be emulated digitally. The remarkable thing about this process is that it is done without any knowledge of how the system actually works—it simply involves learning the system response (output) to input signals. While the concept might seem rather complex, it is a truly simple one—we can take a vintage analog unit, learn its characteristics, then build a plugin that emulates these characteristics. If all the science is done properly, the plugin will sound very much like the original unit.

Convolution reverbs are now widely known and will be discussed in Chapter 24. *Dynamic convolution* is a fairly new technology, developed by a company called Sintefex Audio. Dynamic convolution is far more complex than the convolution process of reverbs. In order to learn the characteristics of compressors (or equalizers), a tedious series of measurements has to be taken before any simulation can be crystallized. Theoretically, it involves studying the compressor response to any possible combination of input levels, frequencies, and settings on the unit itself. In practice, not every possible combination is studied, but a large-scale step-wise approach is taken instead. For example, instead of studying every possible ratio setting, we might only learn integer-based ratios (1:1, 2:1, 3:1, 4:1, 5:1, and so forth). In theory, the emulating processor would only offer the settings that were studied—if only integer-based ratios were learned, the emulator would not offer a 2.5:1 ratio. This could be an annoying limitation. Luckily, the current technology allows dynamic convolution compressors to offer settings that were not studied, and smart algorithms produce an approximate result (a 2.5:1 ratio would produce something between 2:1 and 3:1).

Dynamic convolution is mostly known for its usage in the Focusrite Liquid Channel and Liquid Mix. The former is a hardware unit, the latter is a DSP platform for DAWs. The Liquid Mix, shown in Figure 17.45, includes emulations of 40 compressors (and 20 equalizers) that pretty much constitute a list of the most respectable units one can find in a professional studio.

Figure 17.45 The Focusrite *Liquid Mix*.

18 Limiters

A true limiter has one critical task: to ensure that signals do not overshoot the threshold, *no matter what*. Often it is sudden transients and peaks that we limit. Limiters have many applications outside mixing: they are necessary to protect PA systems, to prevent one FM radio station from intruding on the frequency band of an adjacent station, and, in mastering, to help maximize loudness. In mixing, the task of limiting peaks is, when the final mix is printed, essential where signals must not exceed 0 dB before being converted to digital or bounced to integer-based files (or any other digital media).

We already know from Chapter 17 that a compressor can be configured to behave much like a limiter. To make a compressor a true limiter, we must have peak-sensing, hard-knee, and zero attack time. Compression (well, limiting really) of this nature is the most aggressive and obvious that can be achieved with a compressor. Anything but subtle gain reduction can bring about some dreadful results.

A limiter takes a different approach to the limiting task. The idea is still to have no signals overshooting the threshold, but a limiter employs means to soften the drastic effect of hard limiting. A limiter's principal operation is really concerned with what happens below the threshold and how gain reduction can start before the signal overshoots. One way of doing this is by introducing soft-knee that grows progressively to reach a limiting ratio of ∞:1 at the threshold point. Also, a sophisticated attack function can observe levels below the threshold in order to adapt to the dynamics of the input signal. A look-ahead function on software limiters is standard. Some limiters make use of two stages, where the first stage provides more transparent limiting but no overshoot guarantee; then, if any signals manage to overshoot, they are trimmed down by a second, more aggressive stage.

Figure 18.1 shows a typical limiter with fixed threshold. The three pots on the left are standard for this type of limiter: input level, output level, and release. Some limiters offer variable threshold, but often bringing it down results in respective gain being automatically applied, so the eventual result is the same as bringing up the input level. The input or threshold controls are similar to those of a compressor. The output level often denotes the highest possible output level, also known as *ceiling level*. Often 3 dB of mix headroom is reserved for mastering. Also, even if peaks are limited to 0 dBFS, they can exceed the equivalent analog reference during the digital-to-analog conversion. To prevent this, a little headroom (0.2 dB) is still exercised.

353

Figure 18.1 The Universal Audio *Precision Limiter* plugin.

As opposed to compressors, limiters rarely give the operator any control over the attack, hold, ratio, knee, or peak/RMS sensing. User control over these parameters could defeat the limiting task. As elementary as it may sound, limiters *do not* provide a replacement for compressors—they give us far less control over the final outcome. The novice might be tempted to look at a limiter as a shortcut to compression. But limiters and compressors are worlds apart in the control they offer over the final outcome.

Limiters should not be used as a quick replacement for compressors.

We can look at a limiter as a limited mutation of a compressor. As such, they have their own distinctive **effect** that can be appealing and have a place in the mix. They can also follow a compressor in order to **perfect our compression aims**, such as loudening or level balancing. In that sense, they are used for their original task—**limiting peaks**. Occasional clipping can be handled with a limiter, yet regular clipping would normally call for level attenuation.

Finally, as well as being used for peak limiting, limiters can be used as loudness maximizers. As long as the limiter operation is unnoticeable, we can push the level against it to raise the loudness of the source material. Limiters are often judged based on how hard the input level can be pushed before non-musical artifacts appear. The extra loudness that is such a common part of the standard limiter operation is a risk since it makes A/B comparison unfair—a limiter might fool you into perceiving sonic improvement, even if the limiter is not limiting at all. In mixing scenarios, it is worth using the output control in its conventional compressor role as makeup gain.

Limiters can easily fool you into perceiving a sonic improvement.

Track 18.1: Drums Source
The source track to be limited in the following samples.

Track 18.2: Drums Limited
Mind your monitoring levels—this track is louder than the previous one. This is the result of dialing +9 dB of input gain on the limiter. This might appear to be a substantial improvement compared to the previous track.

Track 18.3: Drums Limited Compensated
This track involves the same settings as in the previous track, only this time the post-limiting signal level is attenuated to match the loudness of the source track. Compare this track to

Track 18.1 to see that the difference between the two is not that great. Also, notice how in the limited version the snare and the kick have lost some attack and have been sent backward.

Track 18.4: Drums Limiting Effect
Setting the limiter's input to the maximum of +24 dB produced this distorting sound, which can be used in a creative context. The level of this track is compensated as well.

Plugin: Universal Audio Precision Limiter

19 Gates

Following compressors, gates are perhaps the second most common dynamic range processors in mixing. Gates are also called *noise-gates*—a name that hints at their traditional usage as **noise** eliminators. In the past, tapes and analog equipment were the main contributors to added hiss. Nowadays, digital technology tends to allow much cleaner recordings. Still, in project studios, unwanted noise on a recording might be the outcome of ground loops, a computer fan, or airborne traffic. Noise, however, is not the only thing gates are employed to reduce. The **rumble** produced by the floor tom when the rest of the drum kit is played is often gated. The hi-hats' **spill** on a snare track and the headphone spill on a vocal track are two more examples. In addition to their traditional role, gates are also used for more creative tasks, such as **tightening drums**, **adding punch**, or **applying dynamic movement**.
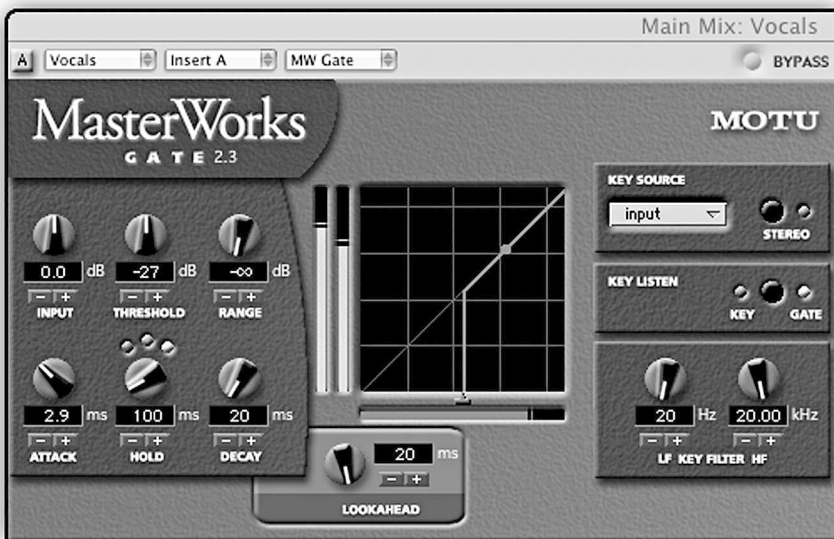
Figure 19.1 A gate plugin. The MOTU *MasterWorks Gate*.
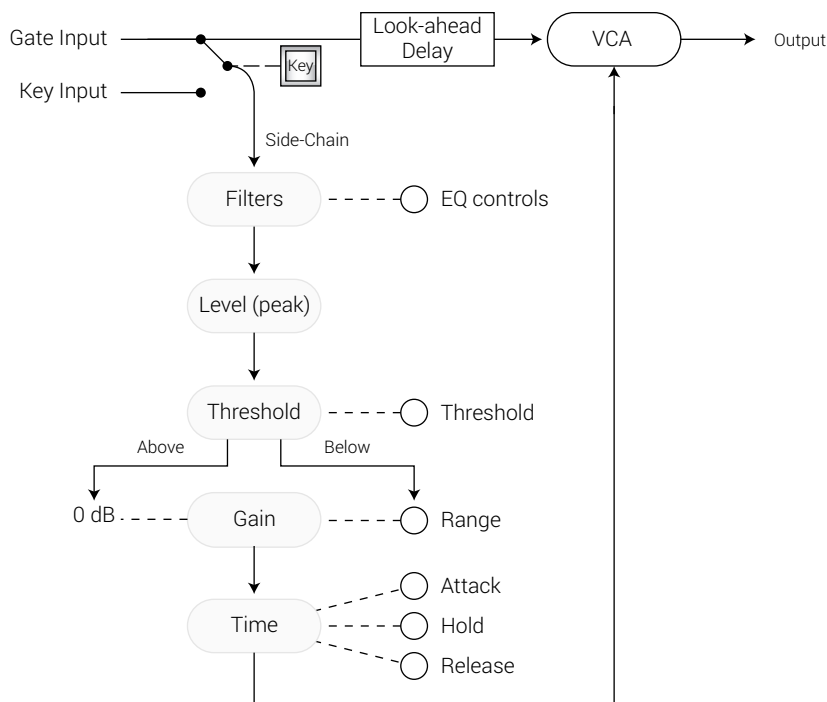
Figure 19.2 Inside a gate. The vertical chain shows the main side-chain stages, and the controls linked to each stage.

Controls

Threshold

Gates affect signals *below* the threshold—these are attenuated, sometimes to a degree that is considered muting. Signals above the threshold pass through unaffected, unless some attack is applied. A gate only cares whether the signal is above or below the threshold; a gate is said to be *closed* when the signal is below the threshold and *open* when the signal is above it. Figure 19.3 illustrates the function of a gate on a snare.

Threshold settings on a gate might seem straightforward—we set the threshold below the wanted signals and above the unwanted signals. In practice, this affair can be tricky since a certain threshold does not always meet both criteria. For example, the snare in Figure 19.3 lost some of its attack and much of its decay. We could retain more of both by lowering the threshold, but this would cause false triggering by the hi-hats. Solutions to this common issue are discussed later on. In the meantime, we should note:

Lower-threshold settings on a gate are often sought as they help in retaining more of the natural attack and decay.
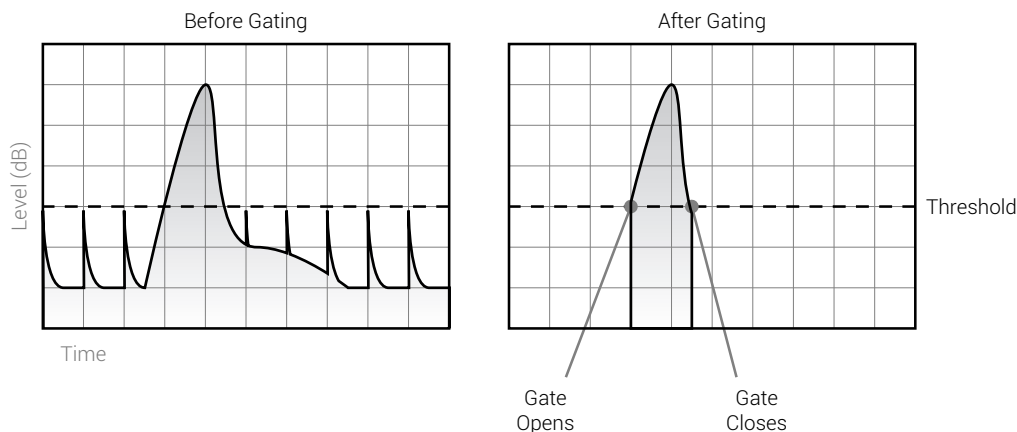
Figure 19.3 Gate threshold function. The original signal before gating involves a single snare hit and a few hi-hat spikes we wish to eliminate. The threshold is set above the hi-hat peaks, so the gate would only open once the snare hits. The gate would close as the snare hit drops below the threshold. After gating, only the loud portion of the snare remains.

Track 19.1: Snare Source
The source snare track is used in the following tracks, where the aim is to gate all drums apart from the snare.

Track 19.2: Snare Threshold –40 dB
A –40 dB threshold is too low. The drums hardly drop below the gate's threshold, and when they do the gate produces abrupt level changes.

Track 19.3: Snare Threshold –30 dB
The kick, toms, and later the hi-hats still overshoot the –30 dB threshold. There are even more abrupt drops in level on this track.

Track 19.4: Snare Threshold –20 dB
Apart from a few tom hits, it is only the snare that overshoots the –20 dB threshold.

Track 19.5: Snare Threshold –10 dB
Only the snare triggers the gate opening here. However, due to the high threshold, some of the snare's leading edge, and much of its decay, are gated as well.

Plugin: McDSP Channel G

Hysteresis

To allow quick gate opening once the signal has exceeded the threshold, the level detection on most gates is based on peak-sensing. Level fluctuations are more profound with peak-sensing than with RMS. While fluctuating in level, signals may cross the threshold in both directions many times over a short period of time. This causes rapid opening and closing of the gate, which produces a type of distortion called **chattering** (see Figure 19.4a).

(a)   Single Threshold

Level (dB)

Time | | | | |
                        O   C   O   C   O   C | O - Open C - Close |

(b)   Threshold with Hysteresis

Level (dB)

                                                          (Opening) Threshold

                                                          } Hysteresis

Time | |
                         Open   Close   (Closing Threshold)
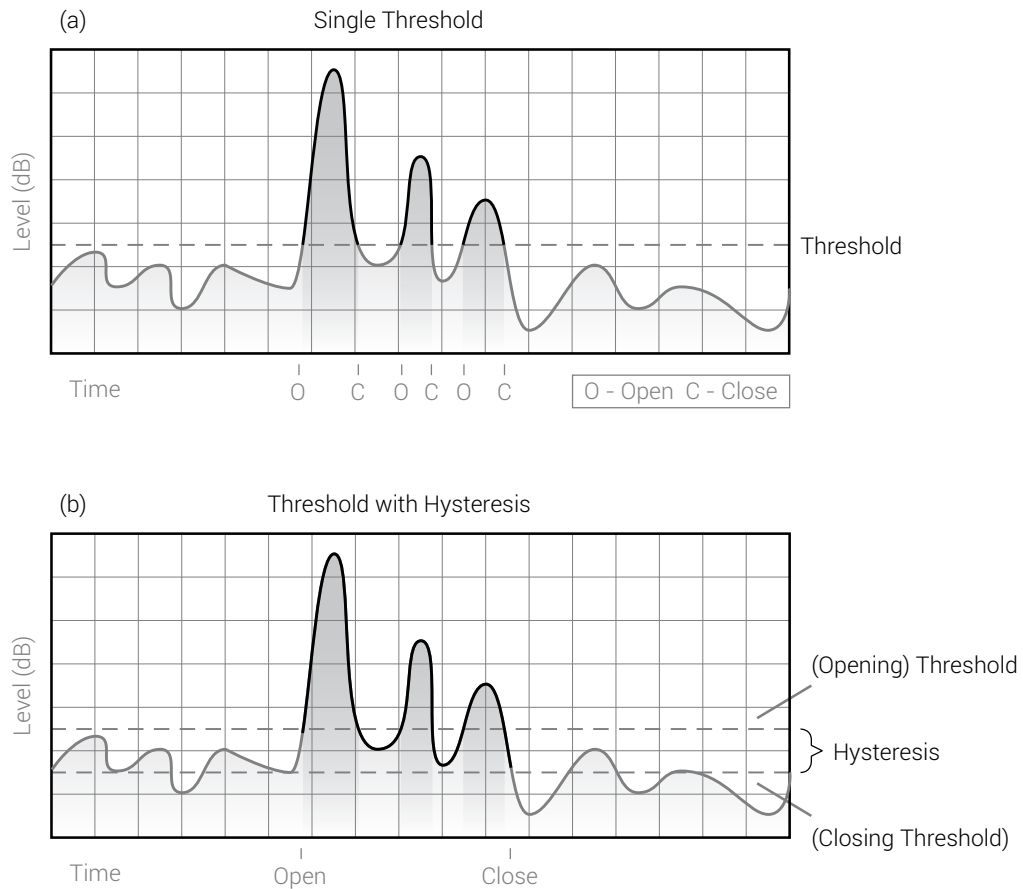
Figure 19.4 Hysteresis control on a gate. (a) A gate with a single threshold and no hysteresis control. Chattering is introduced due to the rapid opening and closing of the gate, which is triggered by quick level changes. (b) A gate with hysteresis control. The gate only opens as the signal overshoots the opening threshold and only closes as the signal drops below the closing threshold. Level variations between the two thresholds do not cause toggling of the gate state; thus, no chattering occurs.

One way to overcome this is by having two thresholds—one above which the gate opens, another below which the gate closes. Having two threshold controls would be cumbersome since adjustments to one will call for adjustments to the other. Instead of providing two controls, gates offer a single threshold and a control called *hysteresis*. The threshold is the opening threshold, and the hysteresis determines how many dBs below it the closing threshold is set. For example, with the threshold set to –20 dB and hysteresis to 5 dB, the closing threshold would be at –25 dB. Figure 19.4b illustrates this.

Many gates do not offer hysteresis as an adjustable control, but have a built-in setting fixed between 4 and 10 dB. When hysteresis control is given, these figures provide a good starting point.

Track 19.6: Kickverb Source
A kick is sent to a reverb, which is gated in the following tracks.

Track 19.7: Kickverb No Hysteresis
The chattering on this track was caused by the reverb level fluctuating around the gate's threshold (–37 dB).

Track 19.8: Kickverb Hysteresis –6 dB
With –6 dB of hysteresis, the gate opens at –37 dB and closes at –43 dB. As can be heard, this reduces chattering but does not eliminate it.

Track 19.9: Kickverb Hysteresis –12 dB
A –12 dB of hysteresis eliminates chattering.

Plugin: Logic Noise Gate
Reverb: Universal Audio Plate 140

Range

Range, or *depth*, defines the amount of gain applied on signals below the threshold. A range of –10 dB means that signals below the threshold are attenuated by 10 dB. Often signals below the threshold are considered muted, although in practice this perceived muting is due to heavy attenuation with the typical range of –80 dB. Figure 19.5 shows the transfer function of a gate, while Figure 19.6 demonstrates the effect of different range settings.
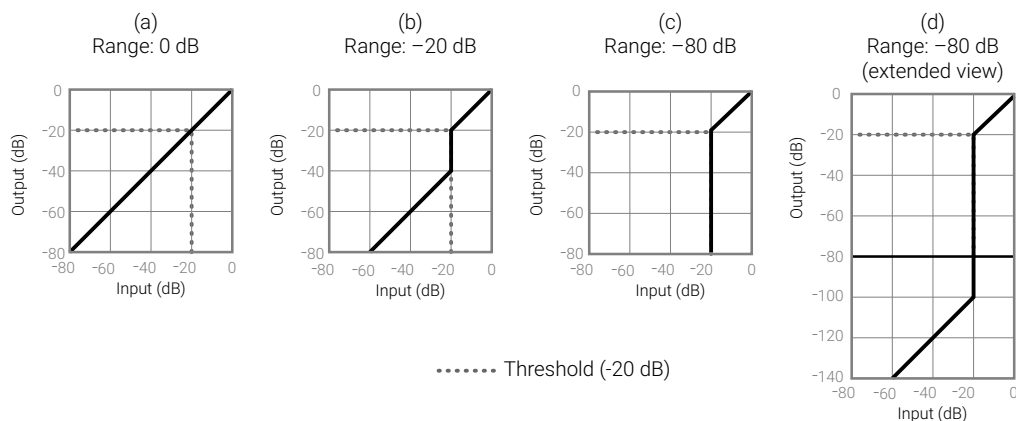
Figure 19.5 The transfer function of a gate. (a) A gate with a 0 dB range: both below and above the threshold the input-output ratio is at unity gain, and the gate would have no effect on the signal. (b) A gate with –20 dB range: all signals below the threshold are simply attenuated by 20 dB. For example, an input signal at –40 dB will leave the gate at –60 dB. (c) A gate with –80 dB range, which effectively mutes all the signals below the threshold. (d) Again, –80 dB range, but with an extended output scale that reveals what happens below the limits of the standard scale. We can see more clearly here that an input signal at –40 dB will leave the gate at –120 dB.
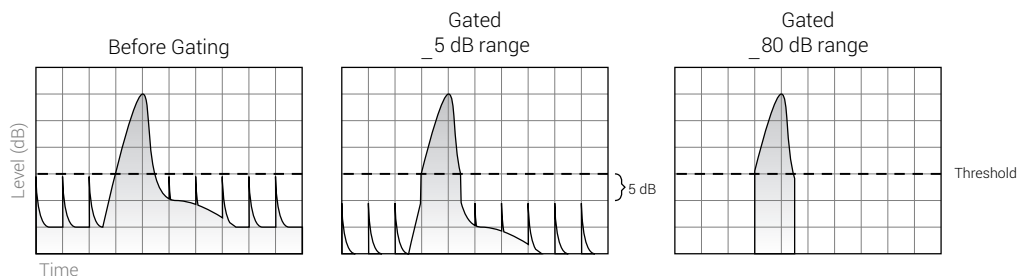
Figure 19.6 The range function of a gate. When the range is set to –5 dB, signals below the threshold are attenuated, but still heard. With large range, such as –80 dB, signals below the threshold become inaudible.

The range value can be expressed in either positive or negative values. Depending on the manufacturer, a gain attenuation of 40 dB might be expressed with a range of 40 or –40 dB. This book uses the negative notation. A small range denotes little effect (e.g., –5 dB), while large range denotes more effect (e.g., –80 dB).

Generally speaking, large range settings are more common in mixing. However, sometimes it is only a gentle attenuation we are after. One example would be gating vocals to reduce breath noises between phrases—removing breaths altogether can sound artificial, so these might only be attenuated. Another example involves reducing the room reverb captured on a recording.

The following tracks demonstrate different range settings on a gate. The larger the range, the less gated portions of the sound can be heard. Notice that it is only with a range of –20 dB or less that the gated material is easily audible.

Track 19.10: Snare Range –60 dB

Track 19.11: Snare Range –40 dB

Track 19.12: Snare Range –30 dB

Track 19.13: Snare Range –20 dB

Track 19.14: Snare Range –10 dB

Plugin: McDSP Channel G

Attack and release

Attack controls how quickly the gate opens; release controls how quickly the gate closes. For example, with the range set to –80 dB, a closed gate would apply –80 dB of gain on the input signal. The attack determines how quickly these –80 dB rise to 0 dB when the gate opens, while the release determines how quickly the gain returns to –80 dB when the gate closes.

The response times on a gate are normally set in milliseconds. Attack times usually range between 0.010 ms (10 μs) and 100 ms. Release times are often within the 5–3,000 ms range. As with compressors, both the attack and release times determine the *rate* of gain change (how quickly gain reduction can rise or fall). For instance, a gate might define that response times are referenced to 10 dB of gain change. With an attack of 1 ms and a range of –10 dB, it will take 1 ms for the gate to fully open; but with –80 dB of range, it will take the gate 8 ms to open. The practical outcome of this is that the range affects the overall time it takes a gate to open and close—where appropriate, we can achieve faster attack and release times by dialing a smaller range. In turn, this promotes smaller range settings than a gate might offer—on a busy mix, a range of –40 dB could suffice to make signals below the threshold inaudible.

It is worth noting that both the attack and release controls have the opposite effect on dynamic envelopes to compressors. A longer attack on a gate means that less of the natural attack is retained; a longer release on a gate means that *more* of the natural decay is retained. Figure 19.7 illustrates this on a gated snare.

On a gate: longer attack, less of the natural attack;
longer release, more of the natural decay.
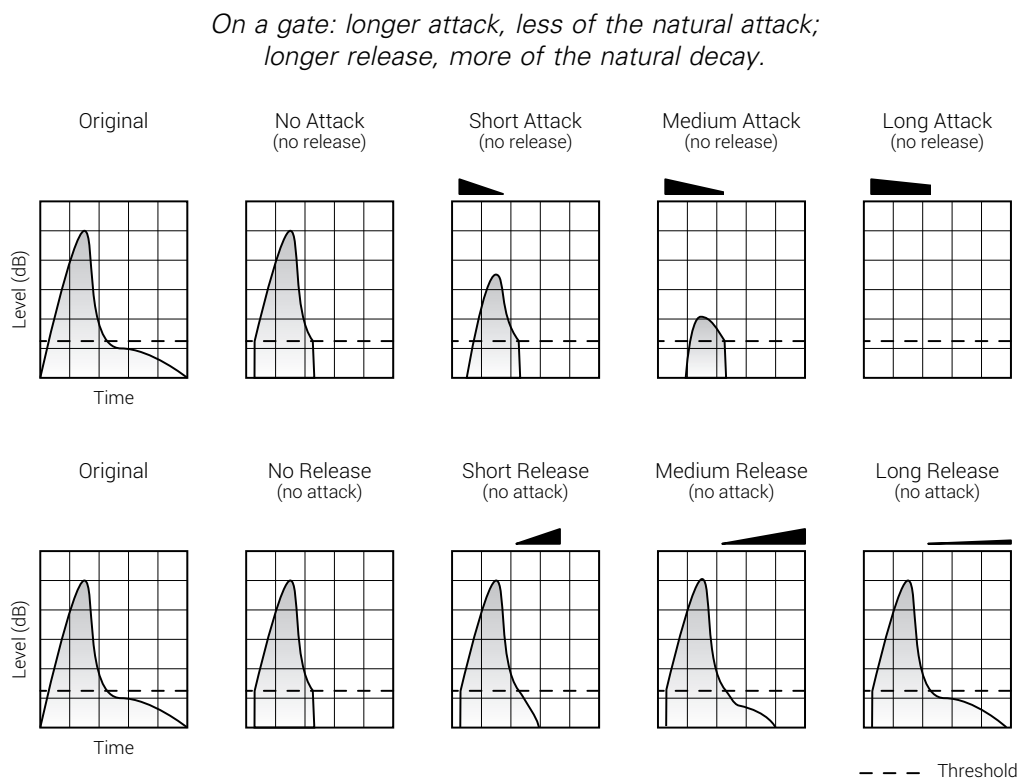
Figure 19.7 The effect of attack and release on a gated snare. The top row shows different attack times with no release; the drop in gain reduction is shown as a triangle above the graphs. We can see that the longer the attack is, the longer it takes gain reduction to drop and the less of the snare's original attack is retained. A long attack setting causes gain reduction to drop so slowly that none of the signal is audible before the input drops below the threshold. The bottom row shows different release times with no attack; the rise in gain reduction is shown as a triangle above the graphs. We can see that the longer the release is, the slower the gate reduces the signal once it falls below the threshold, so more of the natural decay is audible.

Track 19.15: Snare TC Source
The source track used in the following tracks demonstrates different attack times on a snare. The threshold was set to –20 dB, the range to –60 dB, and the release to its minimum of 100 ms. Notice how the longer the gate's attack is, the less of the snare's natural attack is retained.

Track 19.16: Snare Attack 10 ms

Track 19.17: Snare Attack 1 ms

Track 19.18: Snare Attack 5 ms

Track 19.19: Snare Attack 10 ms

Track 19.20: Snare Attack 50 ms

Track 19.21: Snare Attack 100 ms

Another set of tracks with the same gate settings, only this time with attack fixed to its minimum of 10 µs and varying release times. The longer the release, the more of the natural decay is retained:

Track 19.22: Snare Release 100 ms

Track 19.23: Snare Release 300 ms

Track 19.24: Snare Release 600 ms

Track 19.25: Snare Release 900 ms

Plugin: McDSP Channel G
Snare: Toontrack EZdrummer

Using the same principle discussed in Chapter 17, very short attack and release times can cause audible **clicks** due to abrupt level changes. A steep level rise produced by fast attack **generates high frequencies**; the shorter the attack, the higher the frequency (Figure 17.30 illustrates the reason for this). The same applies to release, although the gate operation on the quiet signals below the threshold tends to be slightly less noticeable. Depending on the gated signal, fast response times might also alter the low and mid-frequencies, which essentially means that a gate can affect any part of the frequency spectrum—a side effect we have to observe.

The attack setting has an extra weight when gating transients, especially those with low-frequency content such as kicks. Most of the kick's character is in its very first cycle, with most of the impact gearing up at the beginning of this cycle. A gate is most likely to reshape this important part of the waveform—either with a short attack that produces sharp level changes or with a longer attack that reshapes the waveform and withholds the kick impact. One way or another, the kick's tonality is very likely to change, with the lows, mids, and highs all likely to suffer. Ideally, we would like the attack to retain the original shape of the signal's waveform, which is often only possible with look-ahead (discussed later). Figure 19.8 illustrates this.
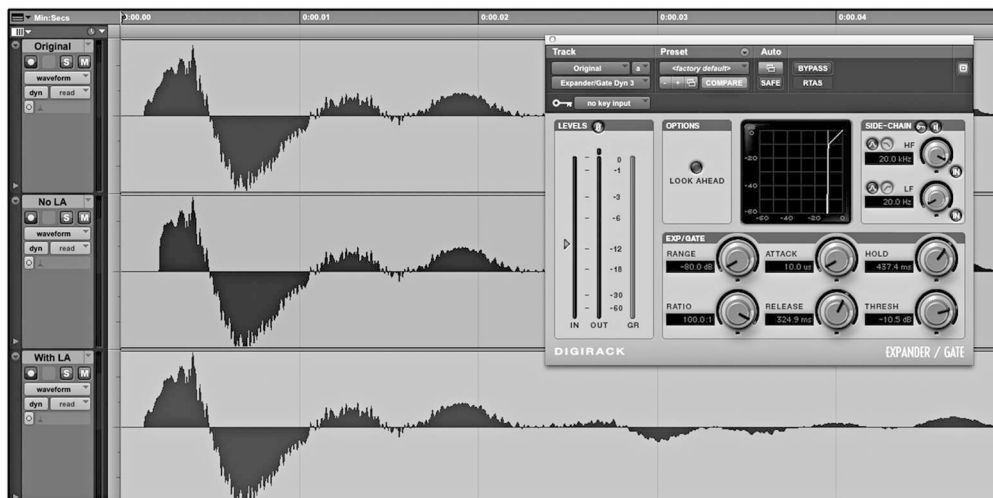
Figure 19.8 A gated kick. The top track shows the original kick, and it is hard not to notice the dominant first cycle. The middle track shows the kick after gating with the settings shown on the Digidesign *DigiRack Expander/ Gate* on the right. The attack was set as short as possible and the look-ahead was disabled. It is possible to see the partial loss of initial impact and the steep level climb caused by the gate opening, which generates high frequencies. The bottom track shows original kick after gating with the same settings, but with look-ahead enabled. You can see how the initial attack was not affected.

Very short settings can also cause **low-frequency distortion** due to the gate operation within the cycles of a low-frequency waveform. Longer attack, release, or hold can rectify this issue. In the case of percussive performance, we must consider how the release and hold settings, which affect the length of the sound, **lock to the rhythmical feel** of the track. In addition, we want the gate to fully close before the next hit arrives, or successive hits will be gated differently (due to variety in gain reduction during the attack phase).

Track 19.26: Kick Source
The source kick track used in the next track.

Track 19.27: Kick Rising Attack
For the purpose of demonstration, the threshold of the gate in this track was set high at –15 dB. The attack was automated to rise from 10 µs to 24 ms. First, notice how the impact of the kick has completely disappeared. Then, notice how the click generated by the steep level change decreases in frequency as the attack is lengthened. Once the attack is long enough, it doesn't generate a click but yields a drop in level (this track was not faded out).

Plugin: Digidesign DigiRack Expander/Gate Dyn 3
Kick: Toontrack EZdrummer

Track 19.28: Kickverb Release Click
The click generated when the gate closes is caused by the steep drop in level—the result of a 0 ms release.

> **Track 19.29: Bass Source**
> This is the source track for the following sample.
>
> **Track 19.30: Bass LF Distortion**
> The extremely unflattering distortion in this track is the result of threshold at –14 dB, range of –60 dB, and 0 ms for all time constants. Although exaggerated in this track, distortion caused by super-fast time constants on a gate is usually as unflattering as this.
>
> Plugin: Logic Noise Gate

One issue with the principal operation of gates is that we often want short attack, so more of the signal above the threshold passes through, and short release, so the signal below the threshold is attenuated quickly. These typical short settings are more obstructive for the many reasons explained above. We have to remember that in many cases, the attack and release are applied on large-scale gain changes, with a range of –60 dB or more. Compressors, conversely, often work on 20 dB or so, and moreover gain changes are not as steep due to the gradual development of the input signal. A gate has no such softening mechanism—it is either open or closed, and there is often quite some gain involved when toggling between the two.

> *The attack and release on a gate tend to be obstructive*
> *due to their typically short settings.*

Short attack and release settings are not always appropriate, though. Longer times are often used when the gated instrument has a long natural attack and release; for example, a synthesized pad that rises and falls in a gradual fashion. Short attack and release will simply trim the parts of the signal that ascend above or descend below the threshold. Long settings will keep some gradual sense for both the natural attack and decay.

Longer attack times might let us lower the threshold by a small amount. The reason for this is that false triggers will not be long enough to become audible. For example, a long attack in Figure 19.3 would allow the threshold to be slightly below the hi-hat peaks. Although the top of each hit will trigger the gate opening, the slow attack will mean the gate will not fully open by the time the hit drops below the threshold, potentially leaving these hits inaudible.

Hold

Once the signal has dropped below the threshold, the hold time determines how long the gate's gain reduction is held unaltered. For example, if 8 dB of gain reduction is applied as the signal undershoots the threshold, a hold period of 2 seconds will mean 2 seconds with 8 dB of gain reduction. Once the hold period is completed, the release phase starts. Gates typically offer a hold period of between 0 and 5 seconds.

Hold often replaces release in the task of **retaining the natural decay** of an instrument. There are two reasons for this. First, we can see from Figure 19.7 that quite a long release setting is needed in order to keep the natural decay. Such a long release is not always practical since it might not end before the next time the signal overshoots
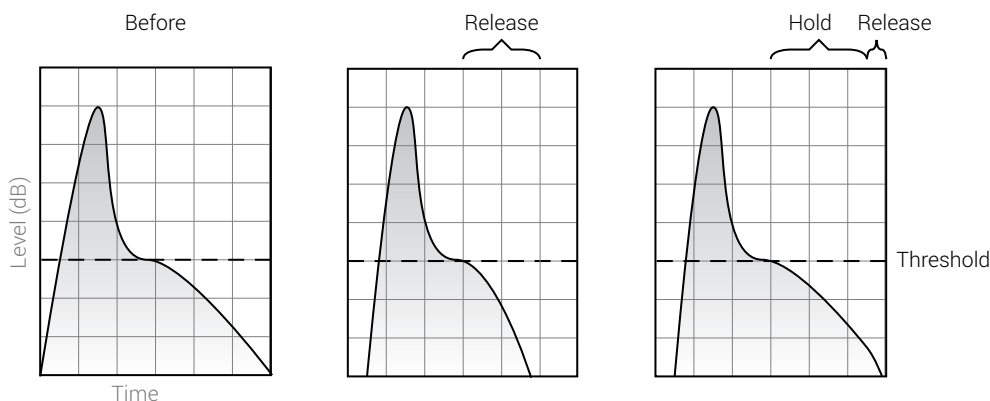
Figure 19.9 The difference between release and hold. With release, the natural decay of the instrument is reshaped, which is not the case with hold. The short release after the hold period is applied to prevent clicks.

the threshold. Second, having the release starting right as the signal drops below the threshold causes an escalated decay—the natural decay of a snare, for example, will be superimposed by the artificial descent caused by the release function. Using hold instead of release lets us keep the fall rate of the natural decay. The hold time can be made to match the interval between two hits, while some release is still used in order for the gate to close without audible clicks (Figure 19.9).

Hold can be used to retain the character of the natural decay.

As with compressors, longer hold periods can reduce **low-frequency distortion**. Longer hold time can also reduce **chattering**, as the gate is held open while the signal rapidly leaps between the two sides of the threshold. In that sense, the hold facility provides a similar function to hysteresis, so when no hysteresis is offered, hold can be used instead. In addition, hold might be used to **compensate for any look-ahead** time that might be in force. For instance, if look-ahead is set to 10 ms, the gate will start closing 10 ms before the signal drops below the threshold. Setting a hold time of 10 ms will compensate for this early response.

The following two tracks demonstrate the use of hold to prevent chattering. The source track is identical to Track 19.7, with no hysteresis. A hold setting of 50 ms is nearly enough to prevent the chattering; 200 ms of hold rectifies it completely.

Track 19.31: Kickverb hold 50 ms

Track 19.32: Kickverb hold 200 ms

Track 19.33: Tom Source
The source track used for the following samples.

Track 19.34: Tom Short Hold Long Release
The gate is set with 50 ms of hold and 710 ms of release. The release reinforces the dropping level of the tom, accelerating its decay.

> **Track 19.35: Tom Long Hold Short Release**
> The gate is set with 710 ms of hold and 50 ms of release. Once the tom drops below the threshold (–17 dB), the hold function keeps the gate open, allowing its decay to pass through unaffected. Then, the release quickly closes the gate. The tom's decay is retained here far better than in the previous track, but at the cost of some extra spill level.
>
> Plugin: Logic Noise Gate

Look-ahead

This is perhaps the ultimate problem with gates: short attack results in clicks; long attack softens the natural attack and can even repress the beginning of words. Lower threshold? Cannot be done or spill will open the gate. Filter the side-chain? Helps a little. But there is a solution—look-ahead.

As with a compressor, look-ahead lets the side-chain examine the input signal slightly before processing takes place. This means that we can dial longer attack times, since the attack phase starts sooner. By the time the signal arrives at the gain stage, gain reduction has decreased to a degree that allows the leading edge of the signal to pass through the gate with no clicks or envelope reshaping. The very same principle applies to release—since the release phase starts shortly before the treated signal drops below the threshold, we can set longer release times. One of the prime benefits of look-ahead is that it allows longer attack and release times, which make the gate operation less obstructive.

Look-ahead makes the gate operation less obstructive.

A look-ahead function on an analog gate will introduce output delay. Software gates have no such issue, provided auto plugin delay compensation is active. Later, we will look at an extremely useful look-ahead trick that results in no delay with both analog and digital gates.

> **Track 19.36 Kick No Look Ahead**
> Using Track 19.26 as a source, the click is caused by a very short attack (10 µs, not automated in this track). The threshold is set to –20 dB and no look-ahead is involved.
>
> **Track 19.37: Kick with Look Ahead**
> By simply enabling look-ahead (2 ms), the click has disappeared and the full impact of the kick is maintained.
>
> Plugin: Digidesign DigiRack Expander/Gate Dyn 3

Side-chain filters

Just like a compressor, some gates let us equalize the side-chain signal that triggers the gating. A classic example of how extremely helpful this is can be seen when gating

a kick track in order to get rid of spillage, such as the hi-hats or snare. It goes without saying that both the hi-hats and snare have less energy on the lows than the kick. So we can filter the highs and mids content from the side-chain, leaving mostly the low-level energy of the kick. In turn, this lets us lower the threshold, so more of the original kick attack can be retained.

Key input

Gates, just like compressors, let us feed an external signal into the side-chain. On a compressor, the external source input is called *external side-chain*. On a gate, the same input is often called *key input* instead (to prevent confusion when a dynamic range processor offers both a compressor and a gate, and the side-chain of each can be fed from a different external source).

We can feed the key input with either a similar signal to the one being gated or a different signal altogether. In the case of the former, an example would involve a *gate micro-phone*—an additional microphone on a specific drum, which is later used in mixdown to feed a gate's key input. A gate microphone has one role—capturing the drum with as little spill as possible. Thus, gate microphones (often dynamic/cardioid) are placed as close as possible to the drum skin and are not concerned with faithfully capturing the timbre of the drum. When we gate a snare, we feed the snare microphone into the gate input and the snare's gate-mic into the key input. This way, we can dial lower threshold, achieve more accurate gating, and have more freedom with the different control settings. Gate microphones aside, sometimes a specific track will benefit from being gated in respect to a similar track. For example, a kick-out track is better gated with respect to the kick-in microphone, which would normally capture less spill (Figure 19.10).
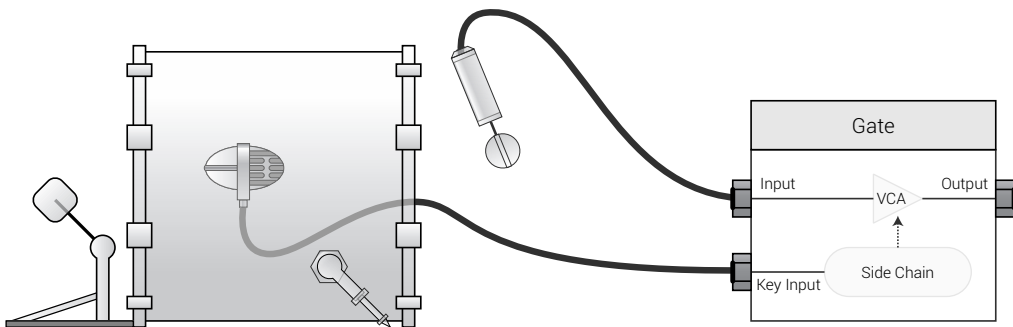
Figure 19.10 Gating kick-out with relation to kick-in. This illustration shows a recording setup, although we usually do this during mixdown with the recorded tracks. The kick-in track is likely to incorporate less spill than the kick-out. While gating the kick-out, feeding the kick-in to the gate's key input is likely to provide more musical gating. Also, the sound arrives at the kick-in microphone slightly before the kick-out microphone, providing a natural look-ahead. A gate-mic can replace the kick-in for other drums such as the snare.

Output level

Although not often provided, the output level control offers a similar functionality to a compressor's makeup gain—boosting or attenuating the entire signal by a set number of dB.

Stereo link

The stereo link function on a gate is similar to that on a compressor—it ensures that both left and right channels are gated identically, so no image shifting occurs.

Meters

Some gates provide a gain reduction meter just like compressors. Some have two indicators—one lights up when the gate is open, one when the gate is closed. Some gates add an additional indicator that lights up during the attack, hold, and release phases, when the gate is neither fully open nor fully closed.

Applications

Removing noise

Noise can be introduced onto a recording in many ways. Tape media, microphones, preamps, any analog component, and A/D converters are just a few examples of systems that superimpose their inherent noise onto a signal. When a recording is carried out in a domestic facility, background noise is also often an issue. Essentially, any signal that ever roamed the analog domain includes some noise. Even purely digital signals, for example a bassline generated with a software instrument, can incorporate either quantization distortion or dither noise.

Luckily, modern technology enables cleaner recordings than in the past; so much of the noise on today's recordings is inaudible. Generally speaking, if high-quality equipment is used, and used correctly, a gate for noise removal might not be needed at any stage of the mix. If, for example, the signal-to-noise ratio on a digital recording is 60 dB (noise level at –60 dBFS), chances are that the noise will not be a problem. Even higher noise levels might not be an issue since in many cases the wanted signal masks it. Noise tends to be more noticeable with sparse arrangements and during the quieter sections of a production—in both cases, there is less to mask it.

One thing worth considering is that noise can become more noticeable as the mix progresses. For example, after applying makeup gain on a compressor, the noise level would rise as well. Also, by boosting the highs of a specific track, the noise is likely to become more noticeable.

Even when the noise is audible, one has to ask: **What is wrong with a little bit of noise?** Many people associate the synthetic sound of digital systems with the lack of noise and distortion. In fact, some engineers deliberately add noise or distortion to a digitally clean recording in order to replicate some of the familiar analog idio-syncrasies. The noise in these cases might be similar to a tape hiss or even crackling vinyl. Nevertheless, other noises such as those generated by washing machines are a problem—they are unlikely to remind anyone of the familiar analog sound.

Some judgment has to be made as to what noise needs gating and what can be tolerated.

Another point is that our ears find varying noise levels (such as breathing) more disturbing than constant noise levels. This fact is taken into account in many noise reduction systems. We have to consider this when gating a noisy track—varying noise levels after gating might be more noticeable than the constant noise level before gating. Say, for example, we have a sparse arrangement with a noisy vocal track. The noise is likely to be masked or be less noticeable while the vocals are sung. The challenge is to make sure that the gate opening and closing does not cause noticeable noise variations. Once the vocal drops below the threshold and diminishes, a slow release can cause a noticeable descent in noise level. If the vocal is sent to a reverb emulator, this drop might be even more noticeable.

Removing spill

Much of today's tracking is done using overdubs; these are less prone to spill, and mostly we deal with headphone leakage and the spill on the various drum tracks. Spill can result in four main problems:

- **Impairing separation**—ideally, we want each track to represent a single instrument or drum (room-mics and overheads are obvious exceptions). A snare track that also contains the hi-hats would make it hard to separate the two. For example, such a hybrid track would restrict independent panning of each drum, and in turn could cause a smeared stereo image on at least one of them.
- **Comb filtering**—the hi-hats on the snare track might not be phase-aligned with the hi-hats track or with the overheads. When the snare is mixed with either track, comb filtering is likely to impair the timbre of the hi-hats and give it a hollow, metallic, or phasing sound. Any instrument might suffer from loss of focus, impact, or timbre coloration if its intended track is mixed with its own spill on another track.
- **Adding dirt**—whenever an instrument is not playing, a spill on its track can produce unwanted sounds. Floor toms are notorious for producing rumble when the rest of the kit is played. Headphone spill might add unwanted noise during quiet sections.
- **Interfering with processing**—to give an example, a loud kick on a snare track might trigger compression and interfere with the snare compression. Brightening a tom track might also emphasize any hi-hat spill it includes.

This list suggests that spill should be removed whenever possible. On occasion, after removing the spill, we may find that bypassing the gate actually has a *positive* effect on the mix. The reasons for this are mostly unpredictable. Regardless, we must remove the spill first in order to learn whether its removal actually improves the mix. It is also worth remembering that processors added later in the mix, such as compressors and equalizers, could also have an effect on this. With that said, once the mix is in its final stages, it is worth trying to bypass spill gates, and see whether the mix changes for better or worse.

One of the main challenges with gating is keeping the timbre of the gated instrument. It has already been mentioned that a lower threshold setting would help in doing so. This task is made harder when the wanted signal and the spill are relatively close in levels, especially if they share the same frequency regions. Snare and hi-hats, kick and toms are potentially problematic pairs, especially if spill was not considered during microphone selection and placement. We may employ any possible gate facility in order to improve our gating. Side-chain equalization lets us attenuate the spill by attenuating its

dominant frequencies, look-ahead lets us retain more of the natural attack, and hysteresis lets us keep more of the instrument's natural decay. For exceptionally tricky gating tasks, it would be extremely beneficial to use at least one of these facilities, if not all of them.

> *One of the main challenges in gating is keeping the timbre of the*
> *instrument, mainly its natural attack and decay.*

When gating drums, there is often a trade-off between the length of the natural decay and the amount of spill—the longer we retain the natural decay, the more spill will escape gating. The spill is often made louder by a compressor following the gate, and the instrument can become louder in the mix for short periods while the gate is open. For example, the hi-hats spill on a snare track might make them louder while the snare decays.

One solution to this involves ducking the intended hi-hats track in an opposite fashion to the snare's gating, so while the open gate adds some hi-hats spill, the actual hi-hats track is attenuated. Another solution is **decay reconstruction**—the gate shortens much of the decay so no spill remains, and a reverb is employed to imitate the missing decay.

Reshaping dynamic envelopes

Just like compressors, gates are also employed to reshape the dynamic envelope of instruments, mainly of percussive ones. We can say that a compressor operates on a transient (once it overshoots the threshold) and on what comes shortly after it (due to the release function). A gate operates on both sides of the transient (essentially, the signal portion below the threshold). Yet, a gate might also affect the transient itself due to the attack function.

As per our discussion in Chapter 17, a part of **adding punch** is achieved by shortening the length of percussive instruments. Typically, the natural decay that we try to shorten is below the threshold, which makes gates a great tool for the job. It is a very old and common practice—you gate a percussive instrument, and with the gate release (and hold) you determine how punchy you want it. Just like with compressors, we must consider the rhythmical feel of the outcome of gating. As gates constrict the length of each hit, the overall result of gating percussive instruments tends to make them **tighter** in appearance. Figure 19.11 demonstrates this practice. The threshold is set to the base of the natural attack and above the natural decay. The range would normally be set moderate to large. The release and hold settings determine how quickly the natural attack is attenuated (with longer settings resulting in longer natural decay). If we only want to **soften the natural decay**, not shorten it altogether, we can use smaller range settings.

> *Gates are often employed to add punch to percussive instruments.*

We can also **accent the natural attack**, **accent transients**, **revive transients**, or **reconstruct lost dynamics** using a gate. This is normally done by setting the threshold somewhere along the transient and dialing a very small range so everything below the threshold is mildly attenuated. It is also possible to boost the output signal, so the
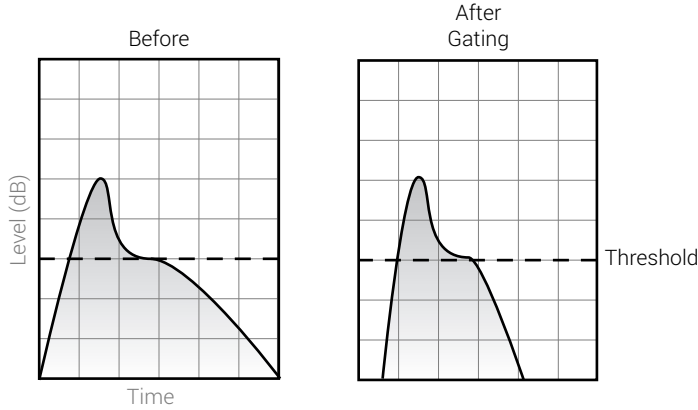
Figure 19.11 Adding punch to a percussive instrument using a gate. The threshold is set to the base of the natural attack and the range is set moderate to large. The release together with hold will determine how quickly the natural decay is attenuated and in turn the overall length of the hit.
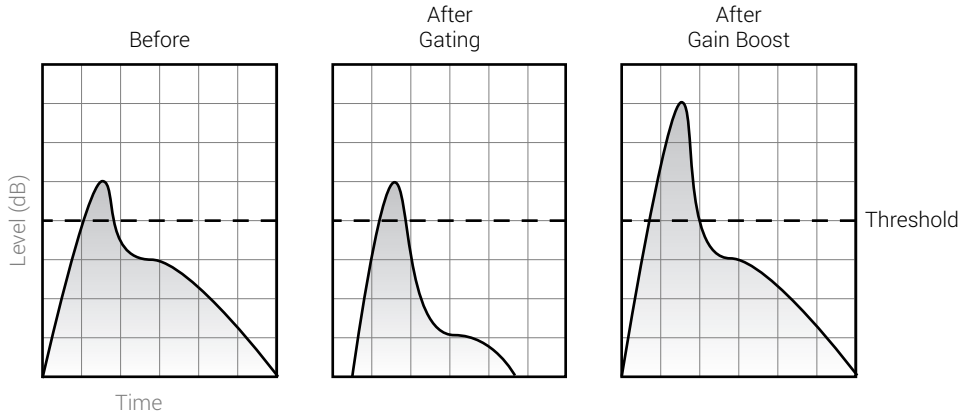
Figure 19.12 A gate emphasizing the attack of a snare. The gate threshold is set halfway through the natural attack. After gating, both sides of the transient have been reduced in level, including the snare's decay. If we boost the gate output, the material below the threshold returns to its original level, but the snare's attack ends up louder in level compared to the signal before gating.

transient is made louder than before but everything below the threshold returns to its original level. Figure 19.12 illustrates this.

It is worth noting that the threshold in Figure 19.12 was set higher than the base of the natural attack. Had the threshold been set lower than that, the gate operation would most likely have altered the dynamic envelope of the snare in a drastic, unwanted way. Figure 19.13 illustrates why this specific application can be somewhat trickier to achieve with a gate than with a compressor. We've already said that a gate's attack and release functions tend to be more obstructive due to their typically short settings. While on a gate we would dial both fast attack and release for this application, on a compressor we would dial longer, less obstructive settings. One outcome of this is that
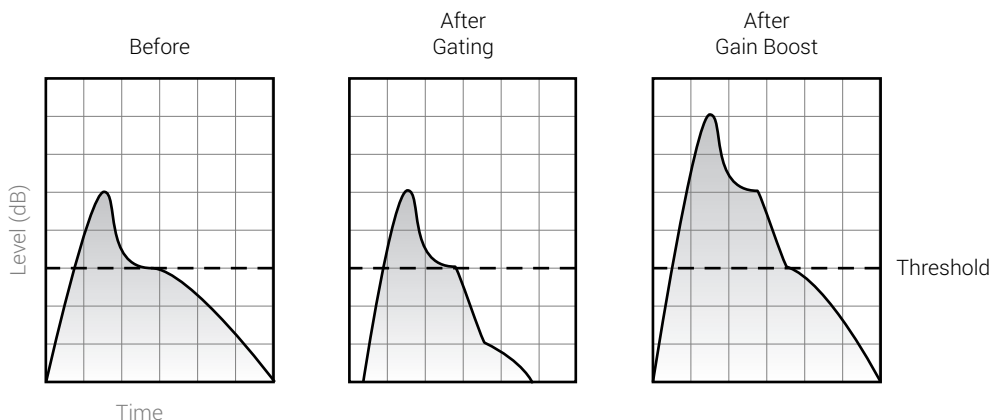
Figure 19.13 A gate reshaping destructively the dynamic envelope of a snare. Due to the low threshold setting, the release function has altered the dynamic envelope in a drastic way, resulting in a new shape that does not resemble the input signal.

when gates are employed for this task, the results tend to be more jumpy than those of a compressor. For these reasons, the gate's range is often kept as small as possible, and look-ahead is employed.

Track 19.38: Kick Snare Source
The source track for the following samples, in which the snare is to be gated.

Track 19.39: Kick Snare Range –4 dB
The gate's threshold was set to –36 dB, the attack to the minimum of 5 µs, hold to 30 ms, and release to 17 ms. The range was set to –4 dB, so the snare's decay is only attenuated a little. Despite no gain compensation, the snare in this track has more profound attack and punch.

Track 19.40: Kick Snare Range –30 dB
With the same settings as in the previous track but –30 dB of range, the snare's decay is removed. This can also translate into a more punchy sound that, despite sounding unnatural, can be suitable for some genres.

Track 19.41: Kick Snare Range Longer Hold
We can lengthen the hold parameter to determine the length of each hit. The settings in this track are identical to those in the previous track, only the hold is lengthened to 90 ms. Compare this track to track 19.38.

The kick in the previous four tracks is also gated, but only in order to get rid of some low-frequency tail:

Track 19.42: Kick No Gate Track

Track 19.43: Kick Gated

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack EZdrummer

In practice

The problem with gates

What is being gated, and how, is largely determined by the threshold setting. Ideally, we would like the gated material to show few gating inconsistencies. This can be achieved if the static threshold setting always affects identical parts of the signal, which is often not the case. For example, if two tom hits vary in level, a gate would attenuate more of the quiet hit's natural attack and decay. Snare hits might be roughly at the same level, but if at a specific section the drummer plays a crescendo, the first few quiet hits might be below the threshold, and therefore muted completely (in this specific case, we can bypass the gate momentarily).

Normally, we loop a specific section while setting the gate. Once happy with our settings, it is important to perform a run-through—playing the track while soloed, beginning to end, to ensure that the gate works as intended throughout the full track length. Even in the case of toms that only play occasionally, it is wise to listen from beginning to end (rather than jumping between the hits) to make sure that no false triggering occurs (for the most part of these run-throughs, we will hopefully enjoy the silence).

Traditionally, gates are placed *before compressors* since the compressor's nature of reducing dynamic range makes gating a trickier affair. However, on a software mixer, it would not be considered odd practice to place a balancing compressor before the gate, so level inconsistencies are reduced before the signal is gated. The gate might still be followed by another compressor that has a different role, such as adding impact.

Gate alternatives: Manual gating

Gating inconsistency is an outcome of varying levels, which are a natural part of a real performance. If we take drums, for example, it could be great having the ability to gate each hit with different settings, but automating a gate for such a purpose would be unwieldy. Fortunately, audio sequencers offer a gating alternative that lets us do exactly this, a practice we will call hereafter *manual gating*. To be sure, manual gating would not produce a characteristic effect like many gates do and can sometimes be senseless—treating 128 kick hits individually can be tedious. But with instruments such as toms, which involve fewer hits, individual treatment for each hit can be feasible. Manual gating can also work on vocals, electric guitars, and many other instruments. It usually takes more time to perform than standard gating, but it gives us more control over the final results.

When sensible, manual gating provides a powerful alternative to standard gating.

Essentially, manual gating is like offline gating that we can then tweak. It involves three main steps, which are demonstrated in Figure 19.14:

- strip silence;
- boundary adjustments; and
- fades.
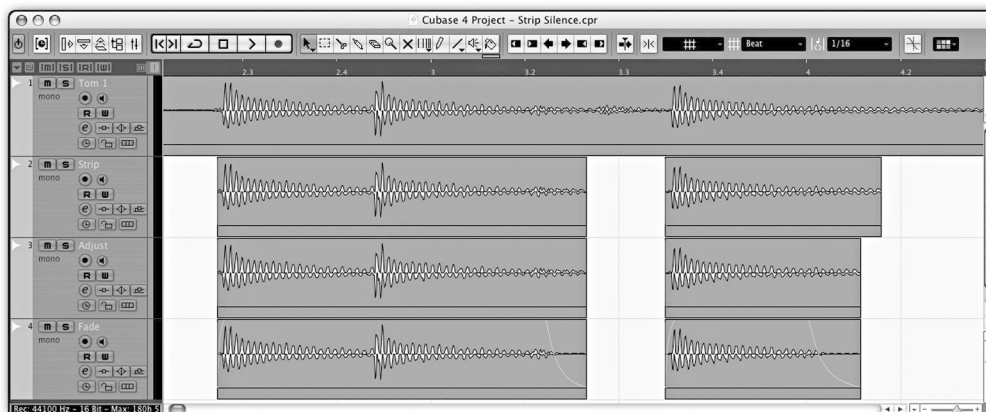
Figure 19.14 Manual gating using strip silence. From top to bottom: the original three tom hits, after strip silence, after boundary adjustments, and after fades. Note that only the right boundary of the second region (third tom hit) and the different fade lengths for each region were adjusted. This screenshot makes use of four tracks to show each step at a time; normally, we perform all steps on the original track.

Strip silence is an offline process that removes quiet sections from audio regions. It works by splicing an audio region into smaller regions that contain the signals we wish to keep. Figure 19.15 shows the "Detect Silence" window of Cubase. Other audio sequencers provide parameters akin to the ones in this figure. The similarity to a gate is evident from the open and close threshold settings. A minimum time for "gate open" and "gate close" is also given, as well as pre- and post-roll times, which create pad margins before and after each splice. The pre-roll time is of great importance; without it, the process is likely to produce clicks, just like a normal gate with a very short attack would.
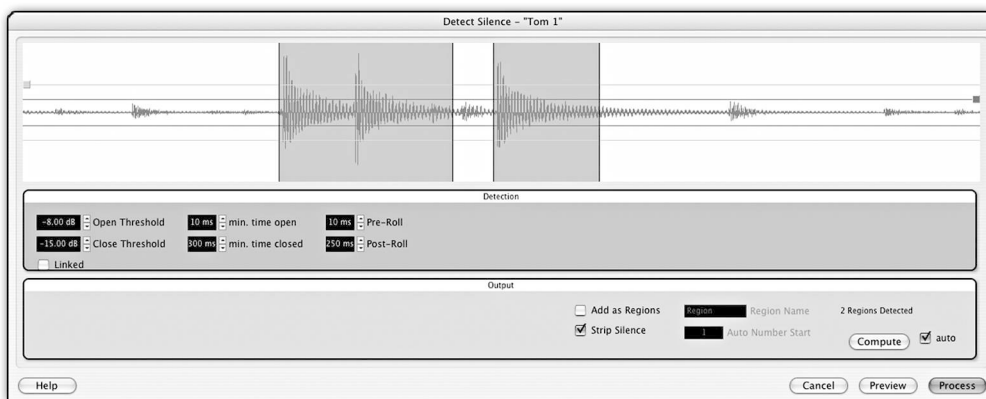
Figure 19.15 Cubase's "Detect Silence" window. In this screenshot, the process is applied on three tom hits. Due to the close proximity of the first two, they end up within the same region. It is also evident that the decay of the third hit is being trimmed (but also some spill that is not visible in this illustration). It is worth remembering that the low-level decay being trimmed here might not be audible unless the tom is played in isolation. One way or another, we can adjust the region boundaries afterwards.

The main aim in the strip silence step is to divide the wanted signal into sensible regions. The exact boundaries of each region are less critical, since the next step involves *adjusting each region's boundaries* individually. If we set the strip silence threshold and pre-roll time appropriately, modifications to the start of each region might not be required. The region's end is adjusted to match the decay length we are after, while taking into account the level and amount of spill. The final step of manual gating involves fading in and out each region, which is similar in effect to the attack and release functions of a normal gate. We can set different fade lengths (attack and release) for each region and normally have control over the shape of each fade.

Gate alternatives: De-noisers

Another alternative to gates, when used to remove hiss or other kinds of noise, is the digital de-noiser (Figure 19.16). Most modern de-noisers need to be taught the nature of the noise to be removed. This is done by feeding pure noise to the de-noiser while in learn mode. Such pure noise is often available at the start and end of tracks. Feeding a de-noiser with anything but pure noise can result in a wired digital effect that might be justified in some creative (somewhat experimental) contexts. Generally, de-noisers do a better job than gates in removing noise and involve a quicker and easier setup. However, this is subject to the noise being of a consistent nature; for example, the static hiss of analog components. De-noisers are not so good at reducing dynamic noise, such as headphone spill or street noise.
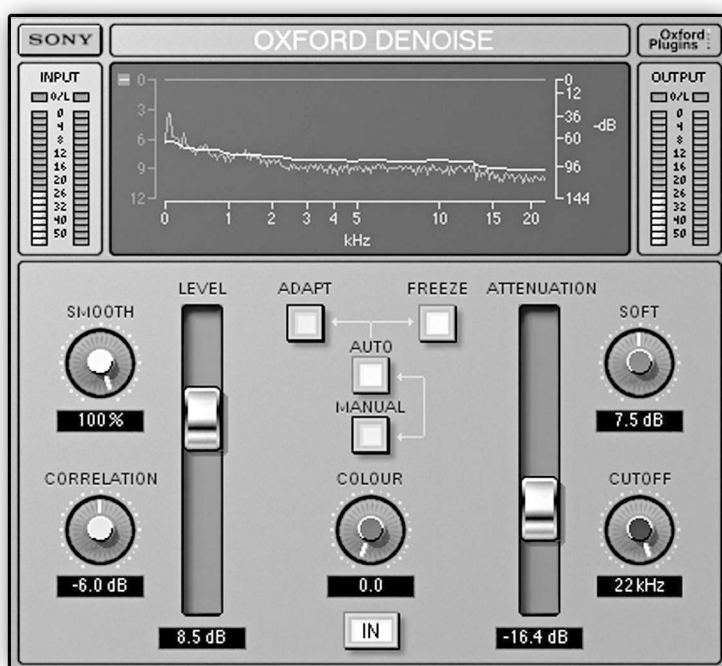
Figure 19.16 A de-noiser plugin, the Sonnox *Oxford Denoise*.

Tricks

Manual look-ahead

A look-ahead facility on a gate is invaluable. Most software gates provide it, but not all hardware units. Often the look-ahead time is fixed, and there might be a switch to bypass this feature altogether. Some gates also let us choose different look-ahead times; for example, the MOTU *MasterWorks* gate in Figure 19.1 (the control below the plot). Whatever gate you use, the trick presented here provides a definitive look-ahead with any gate.

The idea of manual look-ahead is fairly simple. Say we gate a kick, we duplicate the kick track, we nudge the duplicate a few milliseconds back in time, and, instead of routing it to the mix bus, we feed it to the key input of the kick's gate. Figure 19.17 illustrates this arrangement. One advantage of manual look-ahead is that we can set any look-ahead time we fancy. For example, by nudging the duplicate kick 20 ms backward, we get 20 ms of look-ahead time, and we can always nudge the duplicate back and forth (say, in 5 ms steps) to see which time works best. Typically, gates have look-ahead times within the 0–10 ms range, but in specific circumstances even 40 ms could be appropriate. Since the duplicate track is early in time compared with the kick track, the gate's side-chain gets to see the kick a few milliseconds before processing it. The gate, whether software or hardware, has no output delay whatsoever—it is the key input signal being pulled back in time, not the gated kick being delayed.

Note that the *Channel G* plugin in Figure 19.17 is set so the key input signal is filtered (where hi-cut filter is applied to reduce spill). Had the gate not provided side-chain equalization, we could have simply inserted an equalizer on the duplicate track to achieve the same effect.

Figure 19.17 Manual look-ahead in Pro Tools. The top track is the gated kick; the bottom is a duplicate nudged 20 ms backward. The duplicate output is routed to a bus named *Kick SC*, which feeds the side-chain of the McDSP *Channel G*.

Virtually all audio sequencers and digital recorders let us duplicate a track and nudge it backward in time. But nudging is not exactly a part of any tape machine. Therefore, the arrangement would be very similar to the way analog gates implement look-ahead—a track would feed both the gate input and its key input, but the copy sent to gating would be delayed with a delay unit (obviously causing output delay).

While manual look-ahead is extremely useful with gates, it can be used with any other dynamic range processor that has an external side-chain facility—a compressor can benefit from this trick in the same way.

Kick-clicker

It has already been mentioned that short attack settings on a gate can produce an audible click, that the shorter the attack, the higher the by-product frequency, and that the steep level rise can impair the timbre of the gated instrument, especially that of kicks. However, the resultant click caused by fast attack can also have a positive effect on kicks—it adds some definition. Indeed, sometimes gating a kick with a short attack improves its overall presence in the mix, although in most cases we have to compromise with some loss of impact. A simple trick can come to our rescue, enabling us to both keep the impact and add definition.

The arrangement is as follows: the original kick is gated so its leading edge is kept intact (often using look-ahead), while a duplicate track of the kick is gated to produce the click. The two are layered, and we can use the duplicate's fader to determine how much click we want. We can also use the attack and threshold on the clicker-gate to alter the character of the click (instead of using an EQ). Normally, the hold and release are kept very short, so the clicker-gate only opens for a very short period, just to allow the click to pass through. The beauty of this arrangement is that, along with the click, some of the early impact also passes through the gate, resulting not only in improved definition, but also in some added power.

The kick-clicker trick works exceptionally well with sequenced music, where all the hits are of equal or very similar velocity. The problem with a recorded kick is that level variations between hits will produce clicks with changing character. There is an elegant way around this—placing a compressor before the clicker-gate.

Track 19.44: Kick-Clicker
This track is a gated version of Track 19.43. A threshold at –13 dB and the shortest release (5 µs) produce a click.

Track 19.45: Kick Snare with Kick-Clicker
This track is the result of layering the previous track at –9 dB with Track 19.41.

Plugin: Sonnox Oxford Dynamics
Drums: Toontrack EZdrummer

Adding sub-bass to a kick

This trick is used to add some low-energy power, or oomph, to kicks. Although some kicks already contain a hefty low-end, club music and genres such as drum 'n' bass and reggae can all benefit from some extra force, which some describe as chest power, others as LF wind. This trick is also sometimes used in rock mixes.

Figure 19.18 shows the layout of this trick. The idea is to add some low frequency whenever the kick hits. The low frequency is generated by an oscillator, which is gated in relation to the kick—whenever the kick hits, the gate opens to let the oscillator signal through. Figure 19.19 shows this setup in Logic.

Since it is pure power we are after, we do not want the oscillator to produce a recognizable pitch. Thus, a sine wave is used (being the only waveform that generates a pure frequency, with no harmonics to contribute to pitch recognition). The oscillator frequency is often set to 40 or 50 Hz, where 40 Hz would be more powerful, but fewer systems will be able to produce it faithfully. The higher the frequency, the more obvious its pitch becomes, thus frequencies above 50 Hz are rarely used.

To prevent a low-frequency hum that will constantly muddy the mix, the gate is set with maximum range so as to mute the oscillator whenever the kick is not playing. The attack is set short, but it is important to have it long enough to prevent clicks. The clicks produced by a gated low-frequency sine are most likely to vary in nature, as for each hit the gate opens at random points along the waveform cycle (Figure 19.20). A very short release will also produce clicks, which can be even more disturbing as they are often delayed in respect to the kick's attack. With the hold and release, we can set the length of each sub-bass beat. A short beat will simply reinforce the kick, and the longer it is made, the stronger the impact. One creative effect involves making the sub-bass longer than the kick itself, while setting the hold and release to create some noticeable decay (an effect mostly associated with drum 'n' bass tracks). That said, if the sub-bass beat is made too long, we could end up with a constant low-frequency content that will muddy the mix. Another creative approach is to make the sub-bass on the downbeat longer than all other beats, thereby emphasizing the downbeats.

The great challenge with sub-bass beats is sensibly balancing them on the frequency spectrum of the mix. Full-range monitors (or a subwoofer) and a controlled acoustic environment are requisite for this task. The process is usually this: we set the sub-bass level using the fader, and then toggle between the full-range and the near-fields. As we
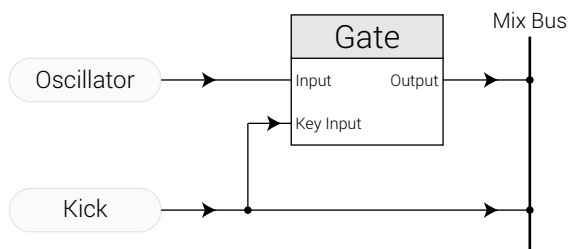
Figure 19.18 Adding sub-bass to a kick. The oscillator is gated in relation to the kick. Every time the kick hits, the gate opens and lets the oscillator signal pass through. Both the kick and the oscillator are mixed.

Figure 19.19 Adding sub-bass to a kick in Logic. The sub-bass track is an aux track on which an oscillator is inserted to generate a 45 Hz sine wave. The oscillator is followed by a gate, which has its side-chain sourced from the kick track (Track 1).

toggle, we would like the relative instrument levels to remain as similar as possible, and, by switching to the full-range speakers, we only want to gain some extra power—it is an extension we are after, not an intrusive addition.

Track 19.46: Kick Snare No Sub-bass
The source track before the addition of sub-bass.

Track 19.47: Kick Snare With Sub-bass
A 50 Hz sub-bass is added to this track. The gate is set with its threshold to –40 dB, range to –80 dB, attack to 4 ms, hold to 150 ms, and release to 400 ms.

Track 19.48: Kick Snare Sub-bass Only
Only the sub-bass from the previous track.

Track 19.49: Sub-bass 30 ms Hold
A slightly tighter effect is achieved by shortening the hold parameter to 30 ms.

> **Track 19.50: Sub-bass 300 ms Hold**
> And a different effect is achieved by lengthening the hold parameter to 300 ms.
>
> **Track 19.51: Sub-bass Varying Clicks**
> This track demonstrates the varying clicks issue. The gate's attack in this track was set to the minimum of 10 μs.
>
> Plugin: Digidesign DigiRack Expander/Gate Dyn 3
> Drums: Toontrack EZdrummer

The arrangement used for this trick is most often employed to add sub-bass to a kick. Another popular version of this trick involves adding filtered white noise to a snare in order to add some definition. We can just as easily add any signal to any other percussive instrument, whether for practical or creative reasons.
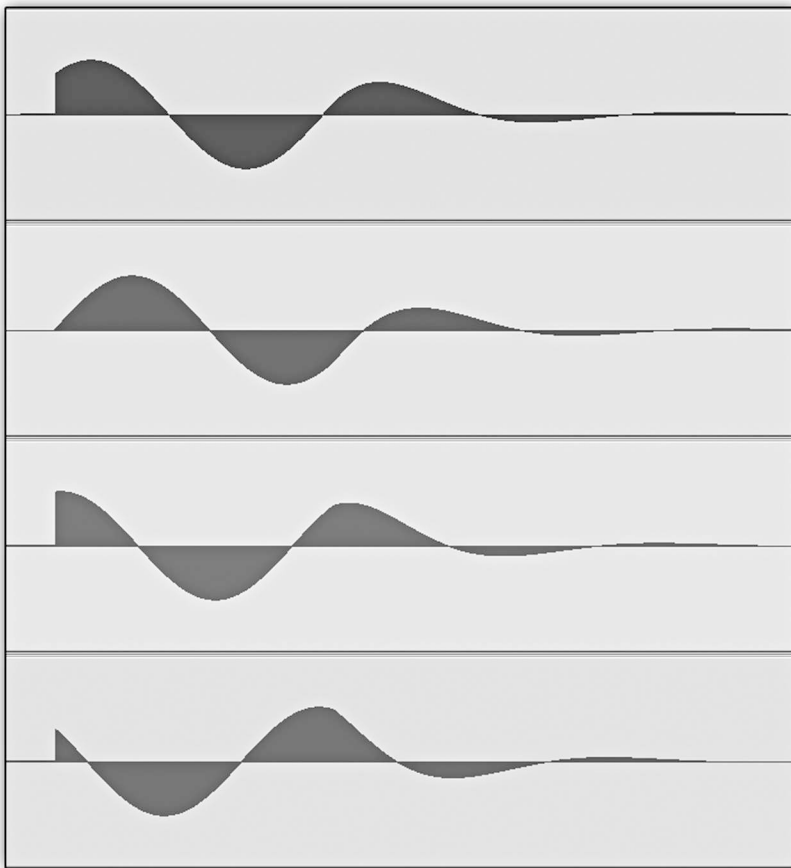
Figure 19.20 Varying clicks on a sub-bass gated with fast attack. A 40 Hz sine wave was gated in relation to a kick. This illustration shows the post-gate outcome of four hits. The short attack produces a click, but, as the gate opens at random points along the cycle, a different click is produced for each hit.

Track 19.52: Kick Snare with White Noise
This track is similar to Track 19.47, only with gated white noise triggered by the snare. The noise was treated with high-shelving attenuation and low-pass filtering. The gate settings are threshold –40 dB, range –80 dB, attack 4 ms, hold 12 ms, and release 260 ms.

Plugin: Digidesign DigiRack Expander/Gate Dyn 3
Drums: Toontrack EZdrummer

Tighten bass to a kick

One way of looking at the previous trick is this: when the kick gets louder, the sub-bass gets louder. We had to mute the sub-bass at all other times to prevent low-frequency hum. How about keeping the same setup, but exchanging the oscillator for a bass guitar or a bassline? The bass will be muted apart from when the kick hits. First, muting the bass between kick hits will result in low-end deficiency (one mixing concept is that the kick provides low-end rhythmical power, while the bass fills the low-end gaps between the hits). Second, we lose important musical information as the bass often provides the harmony root. So this idea is not that great. But how about setting a small range on the gate so the bass is only attenuated by a small amount? The bass guitar will still be audible, but whenever the kick hits it will become slightly louder. This can be very useful if the bass is not extremely tight to the kick. Making the bass louder whenever the kick hits can give the impression that the bassist played the accents tight to the beat.

20 Expanders

Expanders are very common in mixing environments, although not so common in mixes. Expanders are the opposite of compressors in that they expand dynamic range rather than compress it. Indeed, expanders are often used to expand what a compressor has compressed—a do–undo system called a *compander*. Apart from being a vital part of noise-reduction systems for tapes, a compander has little to do with mixing. Still, compressors and expanders work on a very similar principle—both are based around the same design concept and share identical controls, most notably a ratio. The fundamental difference between the two is that a compressor affects signals above the threshold, while an expander affects signals *below* the threshold. If one understands how a compressor works, there is not much to learn about expanders.
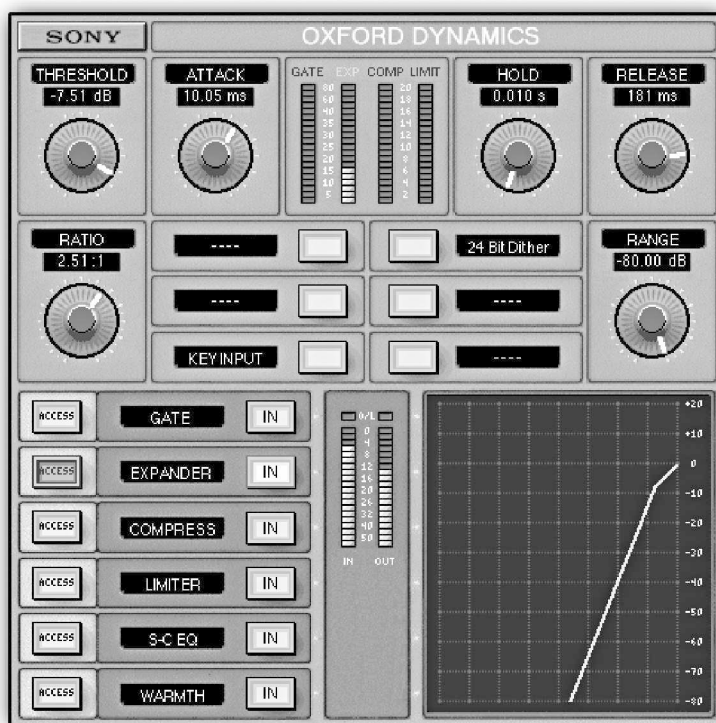
Figure 20.1 An expander plugin. The expander module, part of the Sonnox *Oxford Dynamics*.

Controls

An expander is said to make quiet signals quieter. The ratio, just like with compressors, is expressed with the input-output notation. We can see from Figure 20.2 that, with a 1:2 ratio, a 20 dB fall below the threshold for input signals results in a 40 dB fall for output signals. Visually speaking, the portion of the signal below the threshold is stretched downward to be made quieter. It is interesting to note that with the 1:100 ratio, the transfer function of the expander looks like a gate. Indeed, a gate is often an expander with a high ratio, and most processors are an expander/gate rather than one or the other. Commonly, we have a control that lets us dial ratios between 1:1 and 1:100. On some hardware units, a switch will toggle a gate into an expander with a fixed ratio, often 1:2.
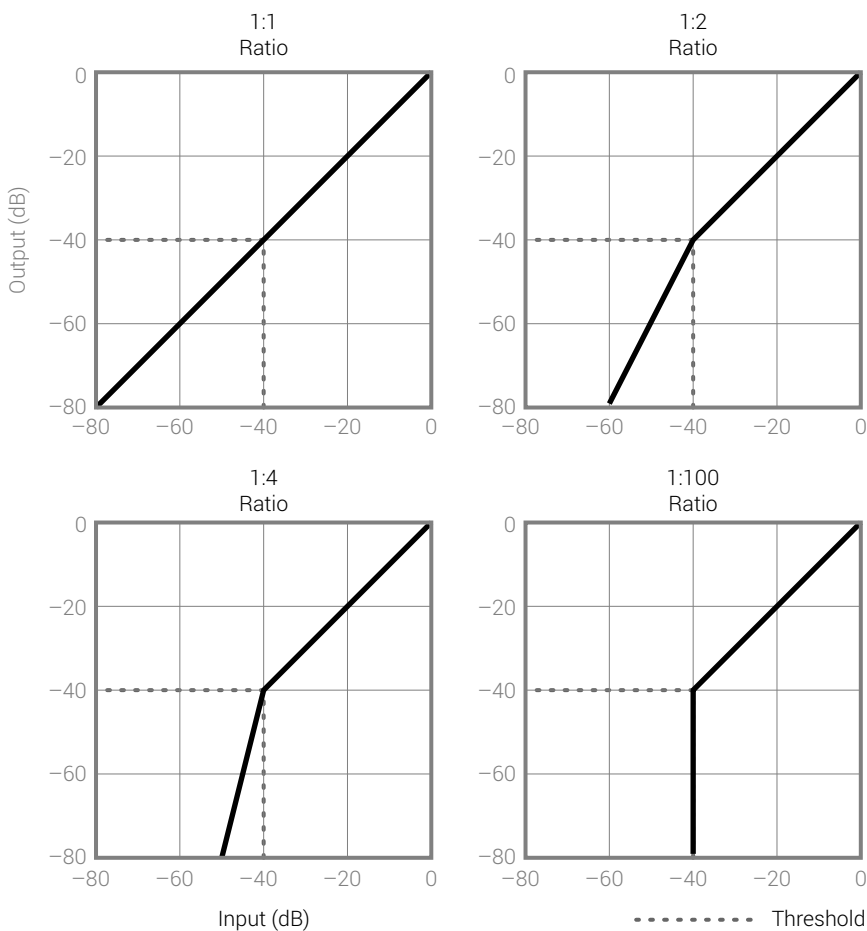
Figure 20.2 The transfer function of an expander with different ratios.

Expander ratios are often given as 2:1 rather than 1:2. This is done for convenience, despite not complying with the standard "input–output" notation. In this book, an expander ratio of 1:8 is higher than an expander ratio of 1:2.

Track 20.1: Piano Source
The source track used in the following three tracks. The sound in this track rises and decays slowly.

The expander in the following track is set with its threshold at –39 dB, the release and attack set to fastest, and the hold set to 1.1 seconds (in order to prevent chattering). Note how the higher the ratio is, the quicker the leading sound rises and the closing sound falls.

Track 20.2: Low Expander Ratio (Piano)
1:2 ratio.

Track 20.3: Moderate Expander Ratio (Piano)
1:4 ratio.

Track 20.4: High Expander Ratio (Piano)
1:16 ratio.

Track 20.5: Snare Source
The source track used in the following three tracks.

The expander's threshold in the following tracks was set to –17 dB, so only the snare would overshoot it. The attack and release were both set to minimum, and the hold was set to 2.9 seconds to maintain some of the snare decay.

Track 20.6: Low Expander Ratio (Snare)
With a 1:2 ratio, the spill is not only audible, but also fluctuates quite noticeably in level. The audible spill is the outcome of insufficient attenuation; the noticeable fluctuations are caused by the gradual transfer curve.

Track 20.7: Moderate Expander Ratio (Snare)
A 1:4 ratio reduces the spill much more successfully, yet some of it can still be heard (although it would probably be masked by the rest of the mix).

Track 20.8: High Expander Ratio (Snare)
The 1:16 ratio in this track yields very fast changes in gain reduction. Compared to the previous track, the function of the expander here is more aggressive, and the spill is reduced in full.

Plugin: Sonnox Oxford Dynamics

Expanders rarely provide a soft-knee option, which means that the transition between treatment and no-treatment at the threshold point can be quite harsh. As with compressors, one way to reduce this artifact is to use the attack and release. In order to understand how the attack and release work, it is worth remembering how a compressor works: the amount of gain reduction is determined by the threshold and ratio settings, and then the attack and release slow down changes in gain reduction. An expander works on exactly the same principle, only the amount of gain reduction is determined by the signal level below the threshold, not above it (the undershoot, not the overshoot amount).

This is worth stressing: both the attack and release on an expander simply slow down changes in gain reduction. More specifically, a gain change from –40 to 0 dB will be slowed down by the attack, and a change from 0 to –40 dB will be slowed down by the

release. Neither the attack nor the release has any relation to the threshold setting. If the signal level fluctuates below the threshold, the attack and release will affect respective changes in gain reduction. Once the signal overshoots the threshold, the applied gain reduction will diminish as fast as the attack allows, and once down to 0 dB the expansion stops. As the signal then drops back below the threshold, expansion starts again, and the release starts to take effect.

The very common hold function on an expander can be implemented in two different ways: the compressor style or the gate style. A compressor-style hold simply alters the release rate, so gain reduction starts very slowly. A gate-style hold freezes the gain reduction for the hold period once the signal has dropped below the threshold.

Most expanders also offer a range control. This parameter defines the maximum amount of gain that will be applied on the signal. For example, a range of –20 dB means that the signal will never be attenuated by more than 20 dB. Figure 20.3 illustrates this.
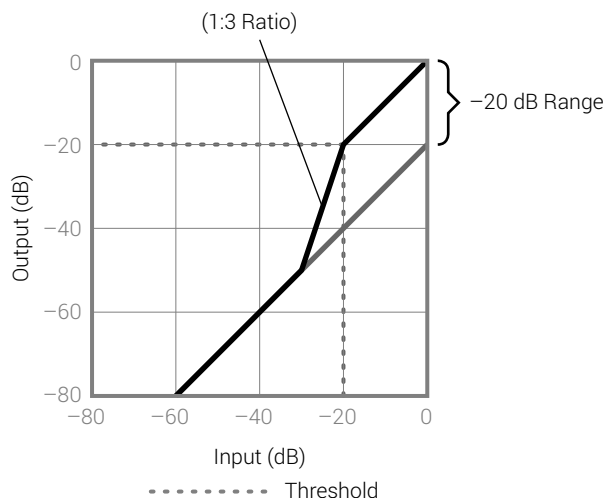
Figure 20.3 Range function on an expander. The gray line denotes the curve of a plain 20 dB attenuation, which sets the bottom limit for the expansion effect. With the resultant ratio curve (black line), the maximum amount of attenuation will never exceed 20 dB.

In practice

While the technical descriptions above might seem somewhat complex, the principal usage of expanders in mixing is very simple—a softer alternative for gates. It was stressed in the previous chapter that gates tend to be obstructive, mostly due to the sharp transition between treatment and no treatment. An expander, with its ratio, provides a more gradual transition between the two. Whether it is large or small attenuation we are after, an expander can give smoother results. The idea is that while a gated signal jumps from one level to another (as the threshold is crossed), an expanded signal slides. Figure 20.4
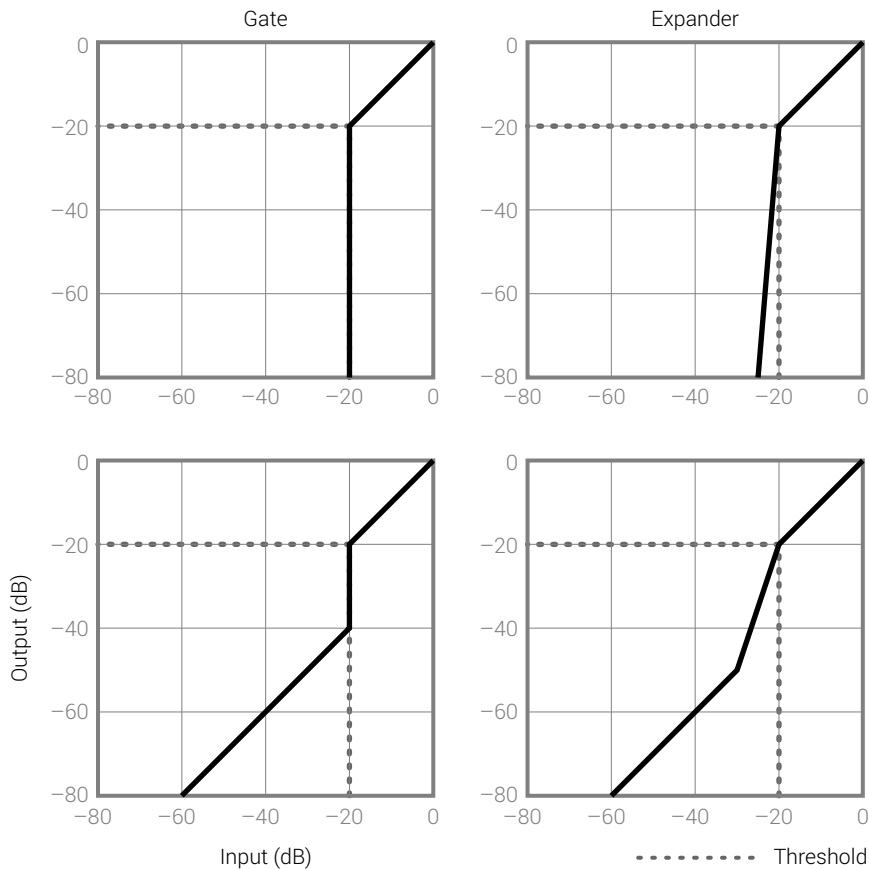
Figure 20.4 Gate versus expander. On the top row, both the gate and expander are employed to drastically attenuate the input signal. While the gate toggles instantly between large attenuation and no attenuation, an expander slides between the two. Even when small range is used, as in the bottom row, the same sliding nature of expanders results in less obstructive effect.

demonstrates these differences. There is an important outcome to this sliding nature of expanders. While on a gate the attack and release are often employed to soften the gate's operation, on an expander this responsibility is consigned to the ratio, allowing the attack and release to concern themselves with the musical aspects of the signal dynamics.

Expanders are used as a softer, less obstructive alternative to gates.

However, there might be some issues with expanders when large ranges are involved. If we assume that large ranges are used when we want to mute signals below the threshold, only gates can ensure such behavior. Expanders can make the very quiet signals inaudible, but signals right below the threshold might still be heard. In order to make these signals less audible, we have to increase the ratio. The issue is that the higher the ratio, the more the expander behaves like a gate—the ratio curve becomes steeper and level changes slide faster. In some situations, signals right below the threshold can only be made inaudible with very high ratios, which effectively turn the expander into a gate. In a way, expanders are more suitable for small range settings.

The following tracks demonstrate how an expander can be used to reduce ambiance from a drum track, and how it differs from a gate.

Track 20.9: Compressed Drums
This is the source drum track from the following samples. The drums are compressed in a way that contains the attack of the kick and snare. Also, the loud ambiance makes the overall drum image appear backward.

Track 20.10: Reduce Ambiance Step 1
The initial expander/gate settings involve a range of –80; a ratio of 1:100 (essentially a gate); and the fastest attack, release, and hold. The threshold was set to –11 dB, which determines what material will not be attenuated by the expander/gate. This material can be heard in this track.

Track 20.11: Reduce Ambiance Step 2
In this step, the range is brought up to –13 dB. Sounds that were lost in the previous track can now be heard. However, like in the previous track, some clicks are evident. Although such treatment can be appropriate sometimes, it might also be considered too drastic.

Track 20.12: Reduce Ambiance Step 3
In this step, the ratio is brought down to 1.5, effectively turning the gate into an expander. The treatment in this track sounds less obstructive than in the previous track. Compare this track to the source track to see how the ambiance is reduced and how the drums' image shifts forward. This can also be described as attack accenting.

Track 20.13: Reduce Ambiance Less Ambiance
We can now control the amount of ambiance with the range control. In this track, it is brought down to –40 dB.

Track 20.14: Reduce Ambiance More Ambiance
This track is the result of –6 dB of range. Compared to Track 20.12, there is more ambiance here.

Plugin: McDSP Channel G
Drums: Toontrack EZdrummer

Another scenario where expanders have an advantage over gates is when we want to de-emphasize the quiet sounds in a gradual fashion. If a gate could talk, it would say: "If it's above my threshold I keep it, if it's below I attenuate it." If an expander could talk, it would say: "Well, if it's below my threshold I attenuate it, and the quieter it is the more I attenuate it." While on a gate, all signals below the threshold are attenuated by the same amount, expanders let us keep more of the louder details (those right below the threshold), and less of the very low details (which might not be heard anyway). This is suitable in situations where it is actually what is above the threshold that we are trying to emphasize, but we do so by attenuating what is below the threshold (which is still important). An example would be adding some dynamic life to a flat drum loop. Figure 20.5 shows a typical transfer function for these types of applications.
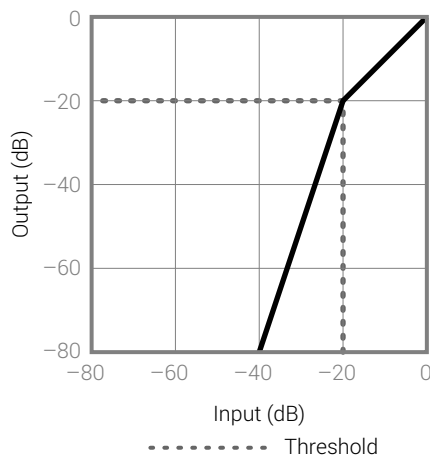
Figure 20.5 Gradual attenuation using an expander. This type of expansion lets us keep more of the loud details below the threshold, and less of the very low level details.

Upward expanders

Why upward expanders are so hard to find is baffling. Both with compressors and gates, we can make the loud signals louder, but we have to take a back-door approach—make what is below the threshold quieter then bring the whole signal up. Upward expanders are designed to make loud signals louder, which makes them an ideal tool for **accenting the natural attack; accenting or reviving transients; reconstructing lost dynamics;** or **adding liveliness, snap, or punch**.

Figure 20.6 shows the transfer function of an upward expander with a 0.5:1 ratio. The rational number denotes upward behavior (some manufacturers will write the same ratio as 1:2). It can be seen from Figure 20.6 that for a 20 dB input rise above the threshold, the output rises by 40 dB. With these specific settings, an input signal entering the expander at 0 dB would come out at +40 dB. Naturally, this would result in hard clipping or overload on most systems. In practice, the ratios used in upward compression are very gentle and do not often go below 0.8:1. Still, if the input signal ever hits 0 dB, the output will exceed the 0 dB limit. Thus, often the output level control is used to bring the overall output level down.

When it comes to upward expansion, we no longer talk about gain reduction, we talk about gain increase. In that sense, the attack controls how quickly the gain increase can rise, while the release determines how quickly it can fall. If a snare is being treated and the threshold is set to capture the attack, shorter attack means more of the natural attack, and shorter release means less of what comes after it. Figure 20.7 demonstrates this.

Both the attack and release functions on upward expanders tend to be far more transparent than on a compressor, since in most cases the expander operates *with* the signal direction, not against it. Despite the dominance of compressors and gates, upward expanders provide an advantageous choice when our goal is *making something*
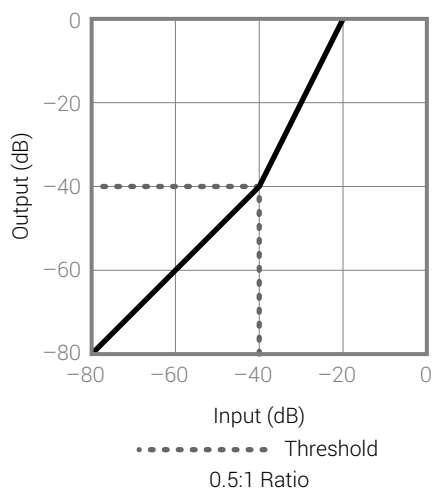
Figure 20.6 An upward expander. Signals below the threshold remain at the same level, while signals above the threshold become louder based on the ratio settings. With the settings shown in this illustration, any input signal above –20 dB will exceed the 0 dB at the output and might overload the system.
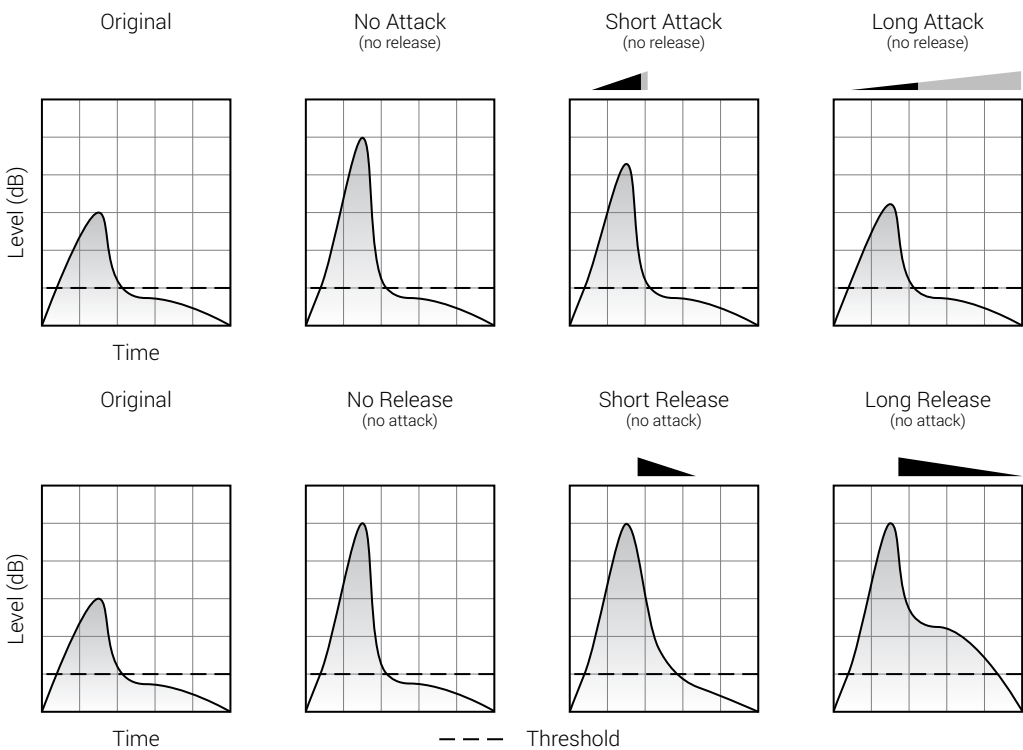
Figure 20.7 The effect of attack and release on upward expansion. The shorter the attack, the quicker the signal above the threshold is boosted, resulting in a stronger effect. Long release can be used to boost the natural decay.

louder. Both compressors and gates are designed to reduce the gain; upward expanders are designed to increase it. If we want to emphasize the impact of a kick, upward expansion is likely to bring natural results, as it simply reinforces the already rising attack.

> *Upward expanders might be the most suitable tool*
> *when it comes to making things louder.*

Track 20.15: Upward Expander
An upward-expanded version of the compressed drums from Track 20.9. The expansion ratio is 0.5:1.

Plugin: Logic Expander
Drums: Toontrack EZdrummer

21 Duckers

We are all familiar with duckers from their common broadcast application—when the DJ speaks, the music is turned down. Most hardware gates provide a switch that turns the gate into a ducker. Even the small dynamics sections on large-format consoles sometimes offer such a feature. For no apparent reason, duckers are hard to find in the software domain. At the time of writing, the *SV-719* in Figure 21.1 is one of the very few plugins that offers ducker functionality. However, ducking as an application is a common part of mixing that is arguably slightly underused.

Figure 21.1 The Sonalksis *SV-719*. This gate/expander plugin also offers a ducker mode.

Operation and controls

Looking back at the internal architecture of a gate in Figure 19.2, all that needs to be done to turn a gate into a ducker is to swap the threshold outputs (above and below). Indeed, duckers offer exactly the same controls as gates and work in a similar way. The sole difference between the two is that while a gate attenuates signals below the threshold, a ducker attenuates signals *above* the threshold. Put another way, once the signal over-shoots the threshold, it is attenuated (ducked). The amount of attenuation is fixed and determined by the **range**. Figure 21.2 illustrates the transfer function of a ducker.
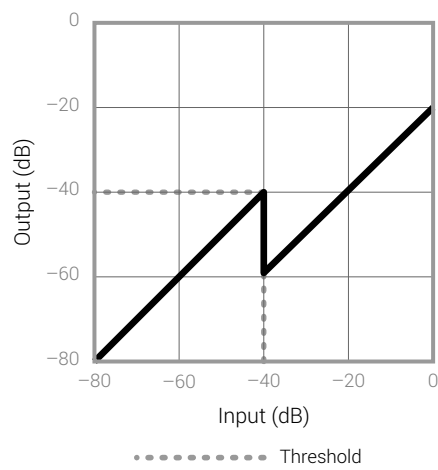
Figure 21.2 The transfer function of a ducker. Signals above the threshold are attenuated by a fixed amount determined by the range (20 dB in this illustration).

The rationale for attenuating a signal by a fixed amount once it exceeds a certain threshold is yet to be discovered. But there are many reasons for attenuating a signal once a *different* signal exceeds a certain threshold, like in the radio example above, where we want the music attenuated when the DJ speaks. To achieve this, we use the key input—the ducked signal (e.g., the music) is fed into the ducker's input, and the ducking signal (e.g., the voice) is fed into the key input. Whenever the voice exceeds the ducker's threshold, the music will be attenuated. Figure 21.3 illustrates this arrangement, and Figure 21.4 shows it in action. Ducking applications always make use of the key input in a ducked/ducking arrangement.

As can be seen in Figure 21.4, gain reduction on the ducked music was applied gradually. This is due to the **attack** and **release** times, which ensure that no clicks occur. In the case of a ducker, the attack determines how quickly the ducked signal is attenuated, while the release determines how quickly the signal returns to its normal level. Also, in Figure 21.4, inspecting the music levels during the attack descent suggests that some music could still interfere with the voice. In addition, there is a hint that once the voice had faded away, the music rise during the release stage was audible.
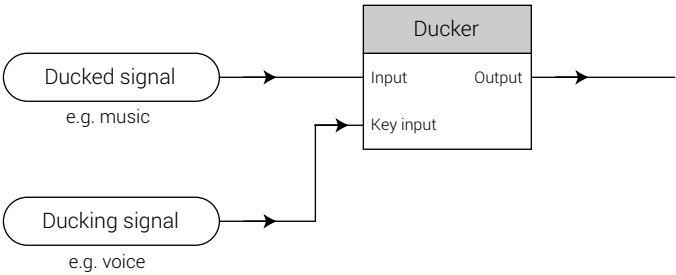
Figure 21.3 The ducked/ducking arrangement. The signal to be ducked is fed into the ducker's input, while the signal to trigger the ducking is fed into the key input. If the music is the ducked signal and the voice is the ducking signal, every time the voice exceeds the ducker's threshold, the music will be attenuated.
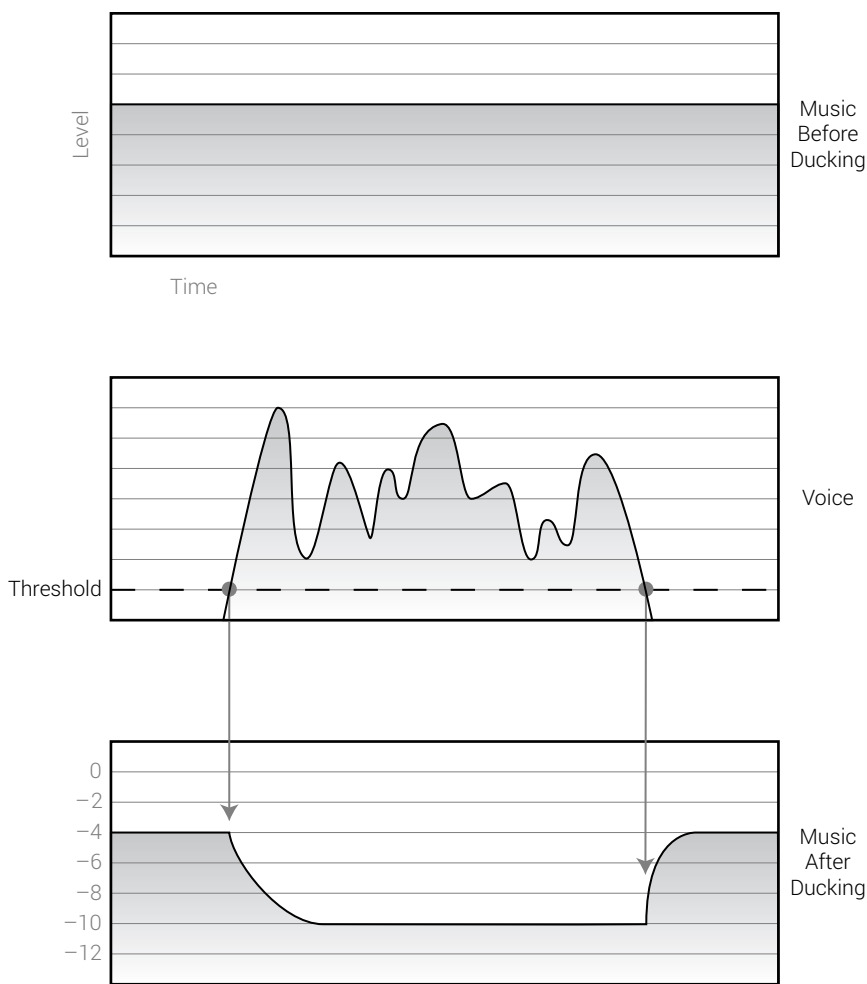
Figure 21.4 Ducking music with respect to voice. The top graph shows the music before ducking, which for simplicity shows no level variations. The middle graph shows the voice that triggers the ducking. The bottom graph shows the music after ducking with a –6 dB range. The arrowed lines denote the points where the voice crosses the threshold, which is where the ducker's gain reduction starts to rise and fall.

Shortening the attack and release might cause clicks and might not be musical. This is where **look-ahead** can be extremely useful—ducking starts and ends slightly beforehand, letting us dial more-musical time constants. The **hold** parameter gives us added control over the ducker behavior once the side-chain signal drops below the threshold.

As already mentioned, often compressors are utilized as duckers. It is worth demonstrating the difference between a real ducker and a compressor used as a ducker. Both processors attenuate the treated signal once the side-chain signal overshoots the threshold. While on a ducker the amount of attenuation is fixed, on a compressor it varies in relation to the overshoot amount. This means that as the side-chain signal fluctuates in level, the amount of gain reduction fluctuates and so does the compressed signal. These differences are illustrated in Figure 21.5.
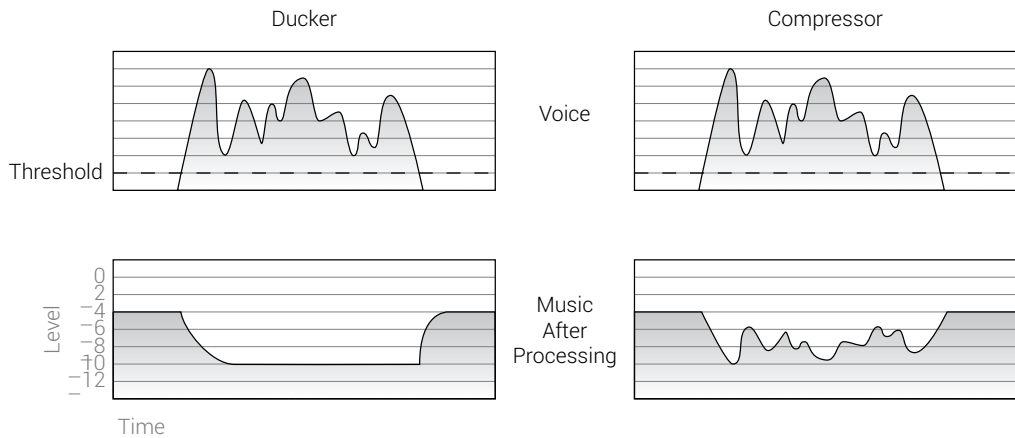
Figure 21.5 A compressor versus a ducker in the ducking challenge. The amount of gain reduction on a ducker is fixed; on a compressor, it is dependent on the overshoot amount. As the side-chain signal fluctuates in level above the threshold, the ducked signal also fluctuates in level. This is not the case with a ducker.

The following tracks demonstrate the difference between a ducker and a compressor-ducker. In both tracks, the uncompressed vocal ducks the guitars. With the ducker, the gain reduction is constant, whereas with the compressor-ducker the gain reduction varies in relation to the level of the voice.

Track 21.1: Ducker

Track 21.2: Compressor Ducker

Plugins: Logic Ducker, Logic Compressor

Applications

In a mixing context, duckers make one instrument quieter when another gets louder—an incredibly useful way to combat **masking**. Say, for example, we want the full slap of the kick but the distorted guitars mask much of it. We can reduce this masking interaction by ducking the distorted guitars whenever the kick hits. It can be surprising how much attenuation can be applied before it becomes noticeable, especially when it comes to normal listeners. Whether done for a very brief moment or for longer periods, ducking is an extremely effective way to clear some space for the really *important* instruments. Indeed, it is mostly the important instruments that serve as ducking triggers and the least important instruments that are the ducking prey. While the options are endless, here are a few common places where this can be done:

- **Vocals ducking their own reverb**—one of the oldest tricks in mixing involves attenuating the vocal reverb whenever the vocalist sings. Vocal reverbs are typically added as an effect rather than for any spatial purposes. In fact, the spatial nature of reverbs can be a great threat to vocals, since it tends to send the vocals back in the

depth field (which is normally not desired). In addition to the ability to mask the vocals, reverbs might blur the intelligibility of the sung words. Ducking the reverb while the vocals are sung makes the vocals more focused and prominent, but we still get the reverb effect during vocal pauses. There are endless accounts of such ducking application in commercial releases.

• **Snares ducking their own reverb**—for the very same reasons as vocals.
• **Kicks and snares ducking their opponents**—it is a typical aim in many contemporary mixes to give an added emphasis to the kick and snare. In dance genres, a potent kick can be one of the prime objectives in the mix. In busy mixes in particular, many instruments can mask the kick and snare, reducing their brunt. By way of analogy, ducking lets the kick and snare push aside any instrument that stands in their way. Both the kick and snare are impulse sounds, so the brief drop they cause on the ducked instruments does not normally result in a great loss of musical information. On some mixes, it is even appropriate to duck the bass with respect to the kick. In addition to kicks and snares, toms can also benefit from the same technique, and, if the majority of the tom sound derives from the close-mics, it is not unthinkable to have them ducking the overheads.
• **A phrasal power guitar ducking a steady power guitar**—some of the heavier rock and metal arrangements involve a power guitar that plays continuously and another guitar that bursts every bar or so. Ducking the steady guitar with the phrasal guitar would cause the steady guitar to drop back (on some mixes even disappear) every time the phrasal guitar has something to say.
• **Ducked effects**—time-based effects occupy some space in the mix over time. Reverbs can cloud the mix, and delays can cause some confusion or draw unnecessary attention. Ducking interfering effects can result in increased clarity and tidiness. Which instrument triggers the ducking is dependent on the arrangement.

> *Remember ducking when you fight masking or clouding, or you are after prominence.*

Track 21.3: Vocal Reverb Normal
A normal arrangement where a vocal is sent to a reverb, but the reverb is not ducked. The dense reverb sends the vocal backward, clouds it, and reduces its intelligibility.

Track 21.4: Vocal Reverb Ducked
The vocal in this track ducks its own reverb. Notice how, during vocal pauses, the reverb gets louder. Compared to the previous track, the vocal here is more focused and forward.

Plugins: Sonnox Oxford Reverb, Sonnox Oxford Dynamics

There is a way to make ducking less intrusive yet more effective, but it involves giving the ducker a rest and using a compressor instead, more specifically a multiband compressor. Back to the earlier example of distorted guitars masking the kick, instead of ducking the entire guitar tracks with the kick, we can feed them into a multiband compressor that is side-chained to the kick. On the multiband compressor, we utilize the high-mids band, so the kick's attack will compress the guitars every time the kick hits. Since no other frequency range will be affected, we can set a deeper ducking to the high-mids of the guitars.

An important goal in mixing is controlled dynamics. Compressors are often employed to contain the level fluctuation of various instruments. However, once levels are contained, the result can be rather vapid—all the instruments play at the same level throughout the song. Duckers can be utilized to insert some **dynamic motion**. Since we have full control over how this is done, we are kept in the zone of *controlled* dynamics. The idea is to add sensible level variations that reflect the action in the mix. Drums are often a good source of ducking triggers; vocals as well. We can, for example, duck the hi-hats in relation to the snare and we can even set a long release on the ducker so the hi-hats rise gradually after each snare hit. Or, we can duck the organ in relation to the vocals or the pad with respect to the lead. This is not done to resolve masking or add prominence; it is simply done to break a static impression and add some dynamism.

Ducking can infuse motion into a dynamically insipid mix.

Then there is always ducking for creative purposes and sheer effect. We can make drum loops and drumbeats far more interesting by incorporating ducking into them. We can feed a delay line into a ducker to create rhythmical level pulses. We can also duck the left channel of a stereo track in relation to the right, and vice versa, to create intentional image shifting.

Track 21.5: Drums No Ducking
This is the source track used in the following track. No ducking is applied on this track.

Track 21.6: Drums with Ducking
The hats and tambourine are ducked by the kick, with relatively long release. Notice how this track has added dynamic motion compared to the previous one.

Plugin: Digidesign DigiRack Compressor/Limiter Dyn 3
Drums: Toontrack EZdrummer

22 Delays

Delay basics

Before the emergence of digital units, engineers used tapes to generate delays (often termed *tape-echo*). Not only is it interesting to see how these delays worked, but it can also be easier to understand the basics of delays using tapes rather than the less straight-forward digital adaptation. Famous units such as the *Echoplex* and the Roland *RE-201 Space Echo* were designed around the principles explained below. Despite many digital tape delay emulations, these vintage units have not yet vanished—many mixing engineers still use them today. Delays are also used to create effects such as chorusing and flanging, which are covered in the next chapter (Figure 22.1).

Delay time

The most basic requirement of a delay unit is to delay the input signal by a set amount of time. This can be achieved using the basic arrangement shown in Figure 22.2. A magnetic tape loop rolls within the unit, where a device called a capstan controls its

Figure 22.1 The Universal Audio Roland *RE-201 Space Echo*. This plugin for the UAD platform emulates the sound of the vintage tape loop machine. The interface looks very similar to that of the original unit.
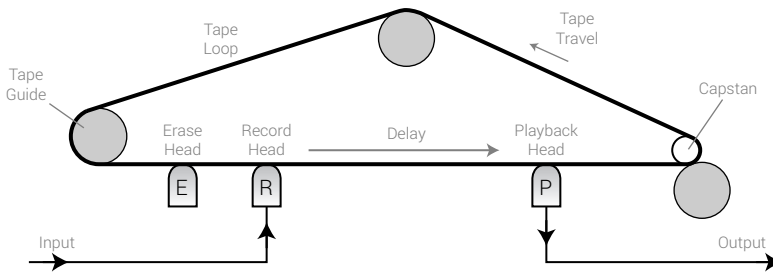
Figure 22.2 Simple tape loop.

speed. The tape passes through three heads: the erase head, the record head, and the playback head (also known as the replay head). After previous content has been erased by the erase head, the input signal is recorded via the record head. Then it takes a period of time for the recorded material to travel to the replay head, which feeds the output of the unit. The time it takes the tape to travel from the record to the replay head determines the delay time.

There are two ways to control the delay time: either changing the distance between the record and replay head or changing the tape speed. Moving any head is a mechanical challenge, and it is much easier to change the tape speed. There is a limitation, however, to how slowly the tape can travel, so some units employ a number of replay heads at different distances, and the rough selection of delay time is made by selecting which head feeds the output. The tape speed is then used for fine adjustments.

For readers completely new to delays, here are a few tracks demonstrating different delay times. The original hi-hats are panned hard-left; the delay is panned hard-right. Notice that both 1 and 25 ms fall into the Haas window, and therefore individual echoes are not perceived.

Track 22.1: HH Delay 1 ms

Track 22.2: HH Delay 25 ms

Track 22.3: HH Delay 50 ms

Track 22.4: HH Delay 75 ms

Track 22.5: HH Delay 100 ms

Track 22.6: HH Delay 250 ms

Track 22.7: HH Delay 500 ms

Track 22.8: HH Delay 1000 ms

Plugin: PSP Lexicon 42

Modulating the delay time

The delay time can be modulated—shortened and lengthened in a cyclic fashion. For example, within each cycle, a nominal delay time of 100 ms might shift to 90 ms, then

back to 100 ms, then to 110 ms, and back to the starting point—100 ms. This pattern will repeat in each cycle. How quickly a cycle is completed is determined by the *modulation rate* (or *modulation speed*). How far below and above the nominal value the delay time reaches is known as the *modulation depth*.

If heads could easily be moved, modulating the delay time would involve moving the replay head closer to or farther away from the record head. But since it is mostly the tape speed that grants the delay time, a special circuit slows down and speeds up the capstan. There is an important effect to this varying tape speed—just as the pitch drops as we slow a vinyl, the delayed signal drops in pitch as the capstan slows down; and similarly, as the capstan speeds up, the delayed signal rises in pitch. These changes in pitch are highly noticeable with both high modulation rate and depth settings.

Track 22.9: ePiano Not Modulated
This track involves an electric piano panned hard-left and its 500 ms delay hard-right. In the following tracks, the higher the depth, the lower and higher in pitch the echo reaches; the higher the rate, the faster the shift from low to high pitch. The depths in the track names are in percentages.

Track 22.10: ePiano Modulated Depth 24 Rate 1 Hz

Track 22.11: ePiano Modulated Depth 100 Rate 1 Hz

Track 22.12: ePiano Modulated Depth 24 Rate 10 Hz

Track 22.13: ePiano Modulated Depth 100 Rate 10 Hz

Plugin: PSP 84

Feedback

Using the arrangement discussed so far, with no feedback applied, the delayed copy would produce a signal echo when mixed with the original signal. It is often our aim to produce repeating echoes, whether in order to simulate the reflections in large spaces or as a creative effect. There is a very easy way to achieve this, which simply involves routing the replay signal back to the record head (Figure 22.3). This creates a feedback loop where each delayed echo returns to the record head, and then is delayed again. Providing a fixed delay time of 100 ms, the echo pattern will involve echoes at 100, 200, 300, 400, 500 ms, and so forth.

A feedback control determines the amount of attenuation (or sometimes even boost) applied on repeating echoes. If this control attenuates, say by 6 dB, each echo will be 6 dB quieter than its preceding echo, and the echoes will diminish over time. If the feedback control is set to 0 dB or higher, the echoes will become progressively louder until the tape saturates (even 0 dB produces such an effect as either tape hiss or the input signal cause increasing level when mixed with the echoes). Such rising echoes are used as an effect, but at some point the feedback
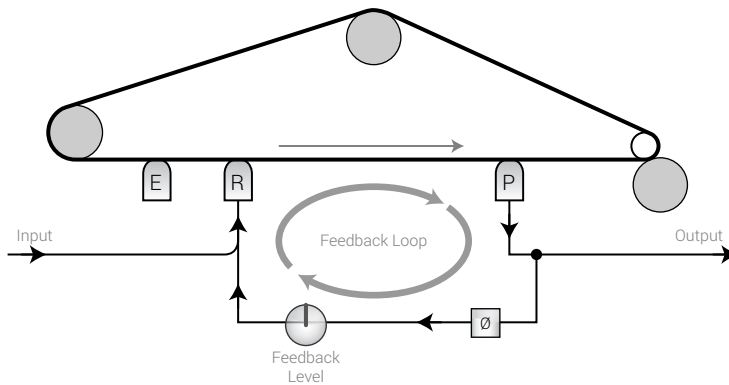
Figure 22.3 Tape delay with feedback loop.

has to be pulled back to attenuation or the echoes will continue forever, causing increasingly distorted sound.

You probably noticed the phase switch before the feedback level control. This switch inverts the phase of the signal passing through the feedback loop (on some units, this is achieved using *negative feedback*). The first echo is phase-inverted compared with the input signal.

But once this echo is delayed again and travels through the loop as the second echo, its phase is inverted again to make it in-phase with the input signal. Essentially, the phase flips per echo, so odd echoes are always phase-inverted compared with the input signal, while even echoes are in-phase. Very short delay times cause comb filtering and can alter the harmonic content of the material quite noticeably. Inverting the phase of signals traveling through the feedback loop alters this comb-filtering interaction—sometimes for the better, sometimes for the worse.

Track 22.14: dGtr No Feedback
This track involves a distortion guitar panned hard-left and 500 ms delay panned hard-right. No feedback on the delay results in a single echo. The following tracks demonstrate varying feedback gains, where essentially the more attenuation occurs, the fewer echoes are heard.

Track 22.15: dGtr Feedback –18 dB
A –18 dB of feedback gain corresponds to 12.5 percent.

Track 22.16: dGtr Feedback –12 dB
A –12 dB of feedback gain corresponds to 25 percent.

Track 22.17: dGtr Feedback –6 dB
A –6 dB of feedback gain corresponds to 50 percent.

Track 22.18: dGtr Feedback 0 dB
With 0 dB gain on the feedback loop (100 percent), the echoes are trapped in the loop and their level remains consistent. This track was faded out manually after 5 seconds.

Track 22.19: Feedback No Phase Inverse
This is the result of a10 ms delay with feedback gain set to –2 dB. This robotic effect is caused by the comb filtering caused by the extremely short echoes.

Track 22.20: Feedback Phase Inverse
By inverting the phase of the signal passing through the feedback loop, the robotic effect is reduced.

Plugin: PSP 84

Wet/dry mix and input/output level controls

Like most other effects, most delays let us set the ratio between the wet and dry signals. There might also be input and output level controls. Figure 22.4 incorporates these controls and shows the final layout of our imaginary tape-echo machine.
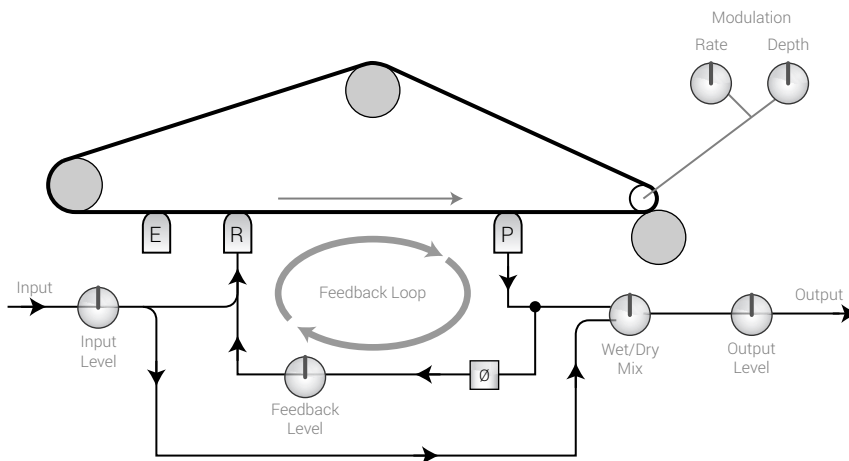
Figure 22.4 A tape-echo machine with all common controls.

The tape as a filter

Tapes do not have a flat frequency response. This is one of the most important characteristics that distinguish them from precise digital media, for recordings in general, and in relation to delays in particular. The most noticeable aspect of this uneven frequency response is the high-frequency roll-off, which becomes more profound at lower tape speeds. Essentially, a tape can be looked at as having an LPF with its cut-off frequency dependent on the tape speed (and quality). In relation to tape delay, each echo experiences some degree of softening of high frequencies. Repeating echoes become progressively darker, which, in our perception, makes them appear slightly farther away.

> **Track 22.21: dGtr Feedback Filtering LPF**
> The arrangement in this track is similar to that in Track 22.17, only with an LPF applied on the feedback loop. For the purpose of demonstration, the filter cut-off frequency is set to 1.5 kHz (which is much lower than what a tape would filter). The result is echoes that get progressively darker.
>
> Plugin: PSP 84

Types

As already mentioned, tape-based delay units are still one of the most revered tools in the mixing arsenal. Around the late 1960s, non-mechanical analog delays emerged, making use of shift register designs, but these were quickly replaced by today's rulers—*digital delays*. A digital delay, often referred to as *digital delay line* or *DDL* for short, works in a way that is strikingly similar to tape delays. Instead of storing the information on magnetic tape, the incoming audio is stored in a memory buffer. Conceptually, this buffer is cyclic, just like the tape is arranged in a loop. There are record and playback memory pointers, very much like the record and playback heads. To achieve smooth modulation of the delay time, sample rate converters are used, providing functionality similar to slowing down and speeding up the tape. Since digital algorithms have an inherently flat frequency response, digital delays often provide an LPF to simulate the high-frequency roll-off typical of tapes. Perhaps one of the greatest limitations of digital delays is that their high memory consumption restricts the maximum delay time. However, we can always cascade two delays to achieve longer times.

Although most digital delays are based on similar principles, there are a few main types.

The simple delay

The simple delay (Figure 22.5) has a single feature—delay time, which is set in either samples or milliseconds. Often this type of process is used for manual plugin delay compensation, but we will see later that even such a simple delay can be beneficial in a mix.

Modulated/feedback delay

Figure 22.6 shows the Digidesign *Mod Delay II* in its mono form. The plugin offers the standard functionality of a DDL: input gain, wet/dry mix, an LPF, a delay time set in milliseconds, modulation depth and rate, and feedback control. The LPF is placed within the feedback loop to simulate the high-frequency roll-off typical of tapes. In the case of the *Mod Delay II*, the phase switch simply inverts the input signal phase, not that of the feedback signal.

The depth parameter is often given in a percentage and determines the delay time alternation range. For example, with a delay time of 600 ms, 100 percent depth denotes 600 ms difference between the shortest and longest delays, which would result in a delay time within the range of 300–900 ms. It is worth mentioning that the

Figure 22.5 A simple delay plugin.
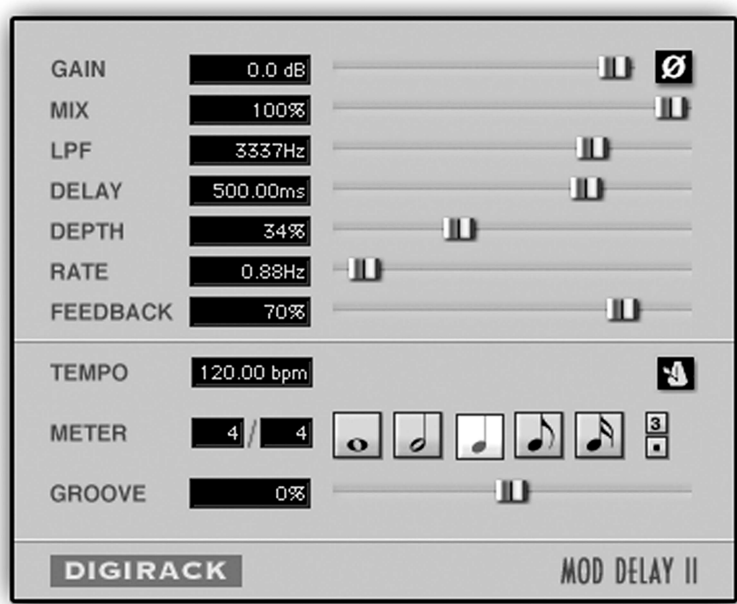
Figure 22.6 The Digidesign *DigiRack Mod Delay II*.

higher the delay time, the more effect the depth will have—while 50 percent depth on 6 ms means 3 ms difference, the same depth on 600 ms would yield 300 ms difference. The modulation rate can be in Hz, but can also be synced to the tempo of the song. There is some link between the modulation depth and rate settings, where the higher the depth is, the slower the rate needs to be for the effect not to sound peculiar.

Figure 22.7 The PSP *84*.

It is worth noting the tempo-related controls in Figure 22.6, which are a very common part of digital delays. In its core, instead of setting the delay time directly in milliseconds, it is set based on a specific note value. In addition to the common values, such as quarter-notes, there might also be triplets and dotted options. Often a dotted quarter-note will appear as *1/4D*, and a half-note triplet will appear as *1/2T*. But the truly useful feature of DDLs is their ability to tempo-sync—given a set note duration, the delay time will alter in relation to the changing tempo of the song. This way, as the song slows down, the delay time is lengthened. On a digital delay unit, this often requires MIDI clock-based synchronization. Plugins, on the other hand, natively sync to the changing tempo of the session. In cases where the musicians did not record to a click, a *tempo-mapping* process can map tempo changes to the tempo ruler.

Appendix C includes delay-time charts for common note values and bpm figures.

From the basic features shown in Figure 22.6, delays can involve many more features and controls. One of the more advanced designs—the PSP *84*—is shown in Figure 22.7. Among its features are separate left and right delay lines, a modulation source with selectable waveforms and relative left/right modulation phases, a resonant filter of three selectable natures, a drive control to simulate the saturation of analog tapes, and even an integral reverb. In addition, the various building blocks can be configured in a modular fashion typical of synthesizers. For example, the modulation section can affect both the buffer-playback speed and the filter cut-off. The comprehensive features on the PSP *84* make it capable of producing a wider variety of time-based effects than just plain delays. But at least some of these features are often also part of other delay lines.

Ping-pong delay

The term *ping-pong delay* describes a delay where echoes bounce between the left and right channels. If both channels' feedback gets sent to the opposite channel, the resultant signal flow would behave the same way as in Figure 22.8.

Although there are some dedicated ping-pong delays (mostly programs within a unit or a plugin), any stereo delay that lets us pan the feedback from one channel to the other can achieve such behavior. Instead of offering feedback pan control per channel, some plugins offer two controls per channel: feedback and feedback crossfeed to the opposite channel. Such a plugin is shown in Figure 22.9.
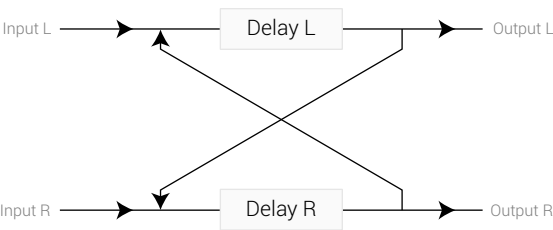
Figure 22.8 A pure ping-pong delay.

Figure 22.9 The Logic *Stereo Delay*.

> **Track 22.22: Ping-Pong Delay**
> This track already appeared in Chapter 14 on panning, but is reproduced here for convenience.
>
> Plugin: PSP 84

Multitap delay

The feedback mechanism, despite its wide use, has some limitations: the echoes are spaced at regular intervals, their level drops in a predictable way, and their frequency content is correlated. We say that each echo is dependent on the echo preceding it. Sometimes we have a very clear idea of the exact number of echoes we would like, what their level should be, or where they should be panned, and we wish to have full control over their spacing and frequency content. A multitap delay is designed to provide such precise control.

Converting a standard tape delay into a multitap delay would involve adding, say, seven playback heads (each considered as a tap), giving us the ability to position each within a set distance from the record head and providing a separate output for each head. These are rather challenging requirements for a mechanical device, but there is no problem

Figure 22.10 The PSP *608* multitap delay.

implementing multitap delays in the digital domain. Figure 22.10 shows the PSP *608* (an eight-tap delay). We have individual control over each echo's timing (horizontal sliders), gain, pan position, feedback send, and filter.

Multitap delay allows precise control over each echo.

Track 22.23: Multitap Source
The source arrangement used in the following tracks, with no delay involved.

Track 22.24: Multitap I
The delay here demonstrates echoes of unequal spacing. Five taps are configured so there is an echo on the second, third, fifth, sixth, and seventh eighth-note divisions. The taps are panned left, right, right, left, right, and the different tap gains make the echo sequence drop then rise again.

Track 22.25: Multitap II
Seven taps are used in this track, arranged in eighth-note intervals. The taps are panned from left to right, and drop in level toward the center. This creates some cyclic motion.

Track 22.26: Multitap III
In this track, seven taps are configured so the echoes open up from the center outwards. An LPF is applied on each tap, with increasing cut-off frequency and resonance per tap.

Plugin: PSP 608
Hi-hats: Toontrack EZdrummer

Phrase sampler mode

This is an added feature on some delay units (also called *freeze looping*) that enables the sampling of a specific phrase and then playing it over and over again. A unit might provide freeze and trigger switches or a single switch that simply enters the mode by replaying part of the currently stored buffer. Many units can tempo-sync this feature so all sampling and triggering lock to the tempo of the song. This feature is different from setting the feedback to 0 dB, since with phrase samplers each repeat is identical to the original sampled material (whereas in feedback mode each echo is a modified version of the previous echo).

In practice

Delay times

Different delay times are perceived differently by our ears and thus play a different role in the mix. We can generalize and say that short delays are used very differently from longer delays. What varies between the two is whether or not we can discern distinct echoes. Yet, we can categorize delay times even further into five main groups. The times shown here are rough and vary depending on the instrument in question:

- **0–20 ms**—very short delays of this sort, provided the dry and wet signals are mixed in mono, produce **comb filtering** and alter the timbre of instruments. Applications for such short delays often involve modulation of the delay time and are discussed in the next chapter. If the dry and wet signals are panned each to a different extreme, the result complies with the **Haas effect**.
- **20–60 ms**—this range of time is often perceived as **doubling**. However, it is the cheapest-sounding approach to doubling, which can be improved by modulating the delay time.
- **60–100 ms**—a delay longer than approximately 60 ms is already perceived by our ears as a distinct echo. Nevertheless, the short time gap between the dry and wet sounds results in what is known as **slapback** delay—quick, successive echoes, often associated with reflections bouncing between parallel walls. Slapback delays are used (in the past more than today) as a creative effect. Since slapback delay could be easily produced using a tape machine, it was extremely popular before digital delays emerged and can be recognized in many mixes from the 1960s.
- **100 ms quarter-note**—such delay times are what most people associate with plain **delay** or **echo**, where echoes are spaced apart enough to be easily recognized. For a tempo of 120 bpm at common time, a quarter-note would be 500 ms. Delay times within this range are very common.
- **Quarter-note and above**—long delay times as such are perceived as **grand-canyon echoes**.

Track 22.27 Vocal Dry
The source track used in the following samples. All the delays in the following tracks involve 30 percent of feedback and are mixed at –6 dB; 414 ms roughly corresponds to a quarter-note, 828 ms to a half-note. Notice that the effect caused by 40 ms of delay can easily be mistaken for a reverb.

> **Track 22.28: Vocal 15 ms Delay**
>
> **Track 22.29: Vocal 40 ms Delay**
>
> **Track 22.30: Vocal 80 ms Delay**
>
> **Track 22.31: Vocal 200 ms Delay**
>
> **Track 22.32: Vocal 414 ms Delay**
>
> **Track 22.33: Vocal 828 ms Delay**
>
> Plugin: Digidesign DigiRack Mod Delay 2

Panning delays

The delay output can be either mono or stereo. In the case of stereo output, where there is a clear difference between the left and right channels, the echoes are often mirrored around the center of the stereo panorama or around the center position of the source instrument. As an example, if vocals are panned center, the echoes might be panned to 10:00 and 14:00. Mono delays can be panned to the same position as the source instrument, but often engineers pan them to a different position, so they do not clash with the source instrument. This not only makes the echoes more defined, but also widens the stereo image of the source instrument. Very often, if the instrument is panned to one side of the panorama, its echoes will be mirrored on the other side. Instruments panned center will normally benefit from a stereo delay, so the effect can be balanced across the sound stage. But there are quite a few examples of an instrument panned center with its mono delay panned only to one side.

> *The following tracks demonstrate various delay panning techniques.*
>
> **Track 22.34: Delay Panning Crossed**
> In this track, the hats are panned hard-right and their mono delay hard-left. The bassline is panned hard-left and its delay hard-right.
>
> **Track 22.35: Delay Panning Crossed Narrowed**
> The hats are panned to around 15:00 and their delay is mirrored to 9:00. The bassline is panned to 9:00 and its delay to 15:00.
>
> **Track 22.36: Delay Panning Outwards**
> Both the hats and the bassline are panned center. The hats' delay is panned hard-left and the bassline's delay hard-right.
>
> **Track 22.37: Delay Panning Mono**
> The hats and their delay are panned hard-left. The bassline and its delay are panned hard-right.
>
> **Track 22.38: Vocal Delay Wide**
> A stereo delay (with no feedback) is panned hard to the extreme. This creates a loose relationship between the echoes and the dry signal.
>
> **Track 22.39: Vocal Delay Narrow**
> With the stereo delay narrowed between 10:00 and 14:00, the echoes are less distinguishable from the dry signal, yet there is a fine sense of stereo width.

> **Track 22.40: Vocal Delay Mono**
> Panning the stereo delay hard-center results in a retro effect that lacks size. Also, the dry signal clouds the echoes.
>
> Plugin: PSP 84
> Hi-hats: Toontrack EZdrummer

Tempo-sync

It would seem very reasonable to sync the echoes to the tempo of the song, especially in the various dance genres. Since delay is the most distinct timing-related effect, offbeat echoes can clash with the rhythm of the song and yield some timing confusion. Delays longer than a quarter-note are especially prone to such errors.

However, there are situations where it actually makes sense to have the echoes out of sync. First, very short delay times, say those shorter than 100 ms, seldom play a rhythmical role or are perceived as having a strict rhythmical sense. Second, for some instruments, tempo-synced echoes might be masked by other rhythmical instruments. For example, if we delay hi-hats playing eighth-notes and sync the delay time, some of the echoes will overlap with the original hits and be masked. Triplets, 3 sixteenth-notes, and other anomalous durations can often prevent these issues while still maintaining some rhythmical sense. Perhaps the advantage of out-of-sync delays is that they attract attention, being contrary to the rhythm of the song. This might work better with occasional delay appearances, where these offbeat echoes create some tension that is later resolved.

It is worth remembering that before audio sequencers became widespread, engineers had to calculate delay times in order to sync to the tempo of the song. Despite the knowledge of how to do this, often delay times were set based on experiment rather than calculations, and, indeed, often the result was not strict note values.

The delay time does not always have to be tempo-synced.

> The following tracks demonstrate various delay times. All apart from Track 22.47 are tempo-synced to a specific note duration:
>
> **Track 22.41: Half-Note Delay**
>
> **Track 22.42: Quarter-Note Delay**
>
> **Track 22.43: Quarter-Note Dotted Delay**
>
> **Track 22.44: Quarter-Note Triplet Delay**
>
> **Track 22.45: Eighth-Note Delay**
>
> **Track 22.46: Eighth-Note Dotted Delay**
> An eighth-note dotted is equal in duration to 3 sixteenth-notes. This type of duration can be very useful since the echoes hardly ever fall on the main beat, yet bear a strong relationship to the tempo.

Track 22.47: 400 ms Delay
Despite having no direct relationship with the tempo of the track (120 bpm), a 400 ms delay does not sound so odd.

Track 22.48: Sines No Delay
The source track used for the following two samples, involving sixteenth-note sequence.

Track 22.49: Sines Tempo Synced
The delay on this track involves a sixteenth-note delay on the left channel and an eighth-note delay on the right channel. Although the delay can be heard, the dry sequence notes and the delay echoes are all tight to sixteenth-note divisions, and therefore the dry notes always obscure the echoes.

Track 22.50: Sines Not Tempo Synced
Having the delays out of sync with the tempo (being 100 ms for the left channel and 300 ms for the right) produces more noticeable delay.

Plugin: Digidesign DigiRack Mod Delay 2
Drums: Toontrack EZdrummer

Filters

Echoes identical to the original source can be quite boring. In addition, if these are quite loud compared with the original, they can be mistaken for additional notes rather than echoes. Although the complexity of tape delay is greater than what simple filtering can achieve, the feedback filters let us incorporate some change and movement into the otherwise statically repeating echoes and distinguish them from the original sound. An LPF makes each echo slightly darker than the previous one, which also makes the echoes progressively undefined and distant. An HPF has the opposite effect—it makes each echo more defined, so, despite their decreasing level, late echoes might be perceived better. Yet they will be distinguishable from the original sound.

Track 22.51: dGtr Feedback Filtering HPF
The arrangement in this track is similar to that in Track 22.21, only with HPF in the feedback loop instead of LPF. The cut-off frequency is 500 Hz.

Track 22.52: dGtr Feedback Filtering BPF
Some delays also offer a band-pass filter. In this track, its center frequency is 1 kHz.

Plugin: PSP 84

Modulation

The modulation feature of a delay line is mostly used with very short delay times. This is done in order to enhance effects such as chorusing or create effects such as flanging. It can be used with longer delay times to add some dimension, movement, or size to a simple delay line, but with anything more than subtle settings the change in pitch becomes evident. This might be appropriate in a creative context, but less so in a practical context.

The automated delay line

One very early and still extremely popular mix practice is automating a delay line. Done mostly on vocals or other lead elements, we often want the delay to catch the very last word or note of a musical phrase. To achieve this, we automate the send level to the delay unit, bringing it up before the word we wish to catch, then down again. Depending on the exact effect we are after, the feedback is set to moderate so the echoes diminish over time, or we set it high to create echoes that barely decline, and then we cut the effect by turning the feedback all the way down.

Track 22.53: Automated Delay Send
Demonstrated here on the vocal, the delay send level was brought up on "I" and then on the closing word *time*.

Plugin: Digidesign DigiRack Mod Delay 2

Applications

Add sense of space and depth

A detailed exploration into our perception of depth is given later in Chapter 24, on reverbs. For now, it is sufficient to say that the very early reflections from nearby walls and other boundaries give our brain much of the information regarding spatial properties and the front/back position of the source instrument in a room. To be frank, reverb emulators are designed to recreate a faithful pattern of such early reflections, and they do so much better than any delay line. However, even the simplest echo pattern caused by a stereo delay can create a vague impression of depth—it is unlikely to sound as natural as what most reverbs produce, but it will still have some effect.

One advantage that delays have over reverbs is that reverbs provide far denser sounds that can consist of thousands of echoes. So many echoes might not only produce an excellent space simulation, but also cloud the mix and cause some masking. Delays, on the other hand, are far more sparse and often only involve a few echoes. When natural simulation is not our prime concern, or in cases where vague space is what we are after, delays may be suitable for the task.

Delays can be used to create a vague sense of space and depth, with a relatively small increase in sound density.

There is another important advantage in the vague simulation that space delays produce—they do not tend to send instruments to the back of the mix as much as reverbs do. This lets us add some sense of space, while still keeping instruments in the front of the mix.

Delays do not tend to send instruments to the back of the mix as much as reverbs.

Compare Track 22.48 with Track 22.49 to see that the sines appear farther back with the delay; the image shifts even farther back in Track 22.50, where the delay is more noticeable. Also, compare Track 22.27 with Tracks 22.29 and 22.30 for a similar demonstration of depth.

Add life

This use, to some extent, is very similar to the previous one—by adding some echoes to a dry instrument, we add some dimension that makes its sound more realistic. Again, the idea that any depth and distance are vague is often an advantage here. Perhaps the main beneficiary of this type of treatment is the vocals—we can add some vitality and elegance, without sending them backward in the mix.

Track 22.54: Vocal with Life
Some listeners will agree that this track sounds slightly more natural when compared to the dry track (Track 22.27). This track involves a two-tap delay mixed at –30 dB.

The comparison between Tracks 22.48 and 22.49 can serve as another demonstration of how the sines—despite being synthesized—sound more natural and "alive" with the delay.

Plugin: PSP MultiDelay 608

Natural distance

If we imagine a large live room where all the musicians are positioned in front of a stereo pair, the sound arriving from musicians right at the back would have to travel farther (ergo for a longer time) than the sound arriving from musicians close to the microphones. As sound takes approximately 1 ms to travel 1 ft, the sound from a drummer placed 25 ft (approximately 8 m) from the microphones would take around 25 ms to travel. It would take around 3 ms for the voice to travel from a singer 1 m away. In many recordings done today using overdubs, there is no such delay—apart from the overhead, all other instruments might be recorded with close-mics, often no farther than 2 ft away. Essentially, when all these overdubs are combined, it would create the impression that all the musicians were playing equidistant from the microphone; in other words, the very same line in space. Although it is primarily reverbs that give us the impression of depth, this very sense can be enhanced if we introduce some short delay on instruments we want farther away in the mix. In this specific case, we only mix the delayed signal without the original one and make no use of feedback or modulation.

Fill time gaps

The echoes produced by a delay can fill empty moments in the mix where nothing much develops otherwise. For this task, a generous feedback is often set along with relatively long delay times, and a filter is used to add some variety to the echoes.

We usually want the delay to fill a very specific gap and diminish once the song moves on, so this application often involves an automated delay line with ridden feedback.

Fill stereo gaps

Delays can solve stereo imbalance problems where one instrument is panned to one side of the panorama with no similar instrument filling the other side. The arrangement is simple—the original instrument is sent to a delay unit, and the wet signal is panned opposite the dry sound. For instruments of a less rhythmical nature, a short delay is probably more suitable. For more rhythmical instruments, say a shaker, the delay time can be longer and tempo-synced. For sheer stereo balance purposes, a single echo (no feedback) would normally do. We can also attenuate the wet signal to make the effect less noticeable.

Pseudo-stereo

It is sometimes our wish to make a mono instrument stereophonic, for example a mono output of an old synthesizer. One way to achieve this involves the same arrangement as above—the mono instrument is panned to one channel, and its delay to the other. It is worth mentioning that this can sometimes also work when applied on one channel of a stereo recording; for example, an acoustic guitar that was recorded with body and neck microphones. If the neck microphone has obvious deficiencies, we can use the body-mic only and create a stereo impression by delaying it and mirroring the delay opposite the dry signal. Applying a filter on the delay would normally make the effect more natural.

Make sounds bigger/wider

Similarly to the filling stereo gaps application above, delays can also be used to make sounds bigger. Often we use a single echo with a short delay time, which is neither too short to cause comb filtering nor too long to be clearly perceived as a distinct echo. This way, we simply stretch the image of a monophonic sound across a wider area on the stereo stage and create an altogether bigger impression. Using this effect may result in loss of focus, but in specific cases it is still a suitable effect.

Let us not forget the Haas trick that was explained in Chapter 12—a common way to enrich and widen sounds by panning a mono copy to one channel and a delayed duplicate to the other channel. The delay time is usually smaller than 30 ms and the effect can give some extra size if the delay time is modulated.

As a key tool in dance music

Delays are one of the most common tools in dance music production. They are used with virtually any combination of settings to enhance many aspects of the production, and for a few good reasons. For one, dance music has a profound rhythmical backbone, and tempo-sync delays can easily enhance rhythmical elements—no other tool has such a strong timing link as delays. Then, sequenced dance productions call for little or no natural sound stage, so delays can easily replace the role of the more natural-sounding reverbs. This doesn't mean that reverbs are not used, but the upbeat tempo

of most dance productions creates a very dense arrangement where reverbs might not have enough space to deliver full effect. In a way, reverbs are more evident in mixes with slower tempos, whereas delays are more evident in mixes with fast tempos. Although we have already established that many recorded music mixes nowadays do not provide a natural sound stage, there are still some limits to how unnatural things can be made before the mix turns purely creative (say, for example, delaying bass guitar or overheads). In dance music, there are hardly any such limitations—delays can be applied on nearly every track, and it takes a while before things start to sound too weird.

23 Other modulation tools

Modulated delay lines are at the heart of a few other effects. Although neither phasers nor tremolos are based on delay lines, their operation is based on modulation. This chapter covers these tools.

Vibrato

Vibrato is the effect caused by rapid and subtle alteration of pitch. Such an effect can be part of a vocal performance, is an extremely common playing technique with string instruments, a known technique on guitars, and can also be the by-product of the Doppler effect as produced by a Leslie cabinet (which was originally called a "Vibratone"). Vibrato can be achieved by modulating a delay line—without mixing the source signal. Subtle settings create similar effects to those that occur in nature. The evident pitch shifting caused by more aggressive settings will make the effect very obvious and unnatural (although this can still be used to creative effect in certain circumstances). The actual settings involve:

- **Dry/wet**—fully wet, dry signal not mixed.
- **Delay time**—between 3 and 8 ms.
- **Modulation depth**—around 25 percent (2 ms). This control is used for more or less effect.
- **Modulation rate**—around 3 Hz.
- **Feedback**—none.

Note that some delay time is still needed. If the delay time is set to 0, it is like having the record and playback heads at the same position. Although mechanically impossible, no matter how much the tape speed was modulated with such an arrangement, the pitch would never alter.

Track 23.1: Original Vocal

Track 23.2: Vibrato (Vocal, Light)
Settings: 100 percent wet, 8 ms delay time, 50 percent modulation depth, 1.5 Hz modulation rate, no feedback.

Track 23.3: Vibrato (Vocal, Deep)
Settings: 100 percent wet, 8 ms delay time, 75 percent modulation depth, 2.2 Hz modulation rate, no feedback.

Track 23.4: Original Guitar

Track 23.5: Vibrato (Guitar)
Settings: 100 percent wet, 8 ms delay time, 40 percent modulation depth, 1 Hz modulation rate, no feedback.

Track 23.6: Original Snare

Track 23.7: Vibrato (Snare)
Settings: 100 percent wet, 8 ms delay time, 65 percent modulation depth, 1 Hz modulation rate, no feedback. In this track, the slow vibrato applied on the snare results in subtle pitch changes. This can be used to make snare samples more realistic.

Plugin: PSP 84
Drums: Toontrack EZdrummer

ADT

Not all singers are fans of the effort involved with real double-tracking. John Lennon was one of them. In search of an artificial alternative, Ken Townshend, then an engineer for The Beatles at Abbey Road Studios in London, came up with a solution. He connected two tape machines, so the delayed vocals from the second were rerecorded onto the first. The modulation on the second machine was achieved using a purpose-installed voltage-controlled oscillator. Having the original voice mixed with its own delayed-modulated version created an effect similar to double-tracking. The technique was termed *artificial double-tracking* (or *automatic double-tracking*) and was used liberally on *Revolver* and subsequent Beatles albums. Townshend's actual arrangement also helped to give birth to chorus and flanging effects.

Although the general consensus is that properly performed double-tracking is better than ADT, having the ability to create this effect artificially during mixdown is a handy one—it gives an impression of fullness, whether by enriching vocals or other instruments. A few artists are associated with the double-tracking effect because they made generous use of it (the late Elliott Smith was one such exponent). But mostly, double-tracking is introduced during specific sections of a song to add some interest and breadth.

The execution of the ADT effect is similar to that of vibrato, except that we mix the dry signal as well:

• **Dry/wet**—possibly at the same level or with dry slightly louder than wet.
• **Delay time**—between 20 and 50 ms.

- **Modulation depth**—40 percent as a starting point.
- **Modulation rate**—less than 1 Hz.
- **Feedback**—none.

> **Track 23.8: ADT (Vocal)**
> Settings: 50 percent wet, 20 ms delay time, 21 percent modulation depth, 0.775 Hz modulation rate, no feedback.
>
> Plugin: PSP 84

Chorus

The actual idea of multiple singers singing in unison was established well before harmony was introduced into Western compositions. Mostly, this was done solely for the purposes of loudness—even today, the size of an orchestra is dependent on the size of the venue and the audience within it. But multiple unison performance is also known to produce an extremely appealing effect called *chorus*. When 20 cellos play the same note, as each player produces a slightly different frequency (within the cents range), and as each cello is either slightly closer to or farther from the listener, the combined effect is mesmerizing. Synthesizers offer a similar feature called *unison*, where the same patch is layered with identical patches slightly detuned.

Perhaps of all the effects mentioned in this chapter, chorus is the most useful for mixing purposes. We might use it in the traditional sense of creating the impression of multiple performance with slightly different pitch and timing. It can, for example, enrich backing vocals and strings. It might also be used to add some richness and polish to plain sounds, or to liven tame organ presets or synthesized strings. Also, it is used to simply make other instruments, such as an acoustic guitar, slightly bigger and wider, and, using more subtle settings, to improve fluidity, on a bass guitar, for example. It can even mask slight pitch detunes in a vocal performance. The denser the mix, the heavier the chorus would need to be to have an effect, but once the settings result in clear pitch alteration, the effect can lose its practical impact and sound gimmicky. In sparse mixes, a touch of chorus will often suffice. One of the qualities of chorus (which can also be regarded as a disadvantage) is that it tends to soften sounds and send them slightly farther away in the mix.

Chorus can be considered a form of ADT with slightly different settings. Chorus tends to sound richer with stereo delays, where slightly different delay times are set on each channel:

- **Dry/wet**—wet set between 20 and 70 percent; the more wet signal, the clearer the effect.
- **Delay time**—between 20 and 80 ms.
- **Modulation depth**—around 10 percent.
- **Modulation rate**—can be anywhere between 0.1 and 10 Hz.
- **Feedback**—none or very little.

Figure 23.1 The Universal Audio *DM-1* in chorus mode. This screenshot shows the settings of the "Chorus Dimension" preset. Essentially, very short delay on both channels, negative feedback of 25 percent, and feedback filter set to 4 kHz; slow rate, little depth, and only 19.7 percent of wet signal, which is phase-inverted.

Track 23.9: Chorus 1 (Vocal)
Settings: 30 percent wet, 21 ms (L), 26 ms (R), 10 percent modulation depth, 0.5 Hz modulation rate, 20 percent feedback. The effect return is panned fully to the extremes.

Track 23.10: Chorus 2 (Vocal)
Settings: 50 percent wet, 40 ms (L), 35 ms (R), 11 percent modulation depth, 1 Hz modulation rate, –25 percent feedback. The effect return is panned fully to the extremes. This track has more noticeable chorus than the subtle chorus in Track 23.9. Note how compared to the original vocal (Track 23.1), the chorused version appears both deeper and wider.

Track 23.11: Chorus 1 (Guitar)
Settings: Same as Chorus 1 (Vocal).

Track 23.12: Chorus 2 (Guitar)
Settings: Same as Chorus 2 (Vocal).

Track 23.13: Original Distorted Guitar

Track 23.14: Chorus 1 (Distorted Guitar)
Settings: Same as Chorus 1 (Vocal).

Track 23.15: Original Bass

Track 23.16: Chorus 3 (Bass)
Settings: Same as Chorus 1 (Vocal), with two exceptions: wet is set to 10 percent, and the effect is panned 50 percent to each extreme.

Plugin: UAD DM-1

Flanging

Just like with guitars, flanging is mostly used as a distinct effect in the mix. Flanging is based on comb filtering created as a result of very short delay times. It yields a series of peaks and dips along the frequency spectrum, which are organized in harmonic series (see Figure 23.2). The actual character of the comb filter is determined by the delay time, the feedback phase switch, and the amount of feedback. The modulation sweeps the comb up and down in frequencies, producing a whooshing sound. Modulation rates are usually kept low and might be synced to a half-note or a one-bar duration. If any waveform option is given, either a sine or a triangle would normally be used.
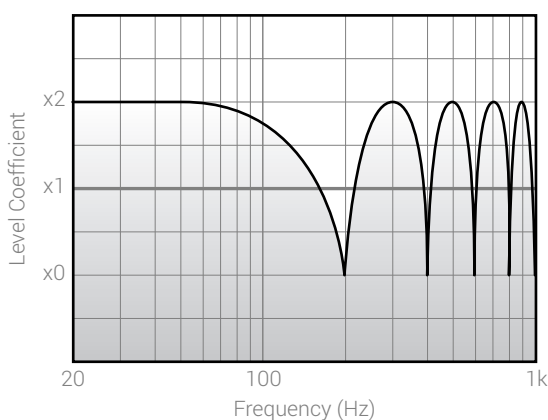
Figure 23.2 The comb-filter response caused by a flanger. The frequencies of the peaks and dips are harmonically related. As the effect is modulated, this response pattern expands and contracts, but the harmonic relationship remains.

By way of analogy, the flanging effect is like a siren where the depth determines how low and high the siren goes, and the rate determines how quickly it goes. The feedback can be thought of as a resonance controller: the higher the feedback, the more the resonance. Phase-inverting the feedback will normally produce a bolder effect:

- **Dry/wet**—50/50.
- **Delay time**—between 1 and 20 ms.
- **Modulation depth**—often higher than 25 percent.
- **Modulation rate**—slow, often tempo-synced to a half note or a bar.
- **Feedback**—the more feedback, the more resonance.

The following tracks were produced using both the light and deep flanger presets of the PSP 42. The deep flanger preset can be seen in Figure 23.3. *The light flanger preset is only different from the deep one in that its feedback is set to 0 percent.*

Track 23.17: Flanger (Vocal, Light)

Track 23.18: Flanger (Distorted Guitar, Light)

Track 23.19: Flanger (Distorted Guitar, Deep)

Track 23.20: Original Drums

Track 23.21: Flanger (Drums, Light)

Track 23.22: Flanger (Drums, Deep)

Track 23.23: Original Shaker

Track 23.24: Flanger (Shaker, Light)

Plugin: PSP Lexicon 42
Drums: Toontrack EZdrummer

Figure 23.3 The PSP *Lexicon 42* in deep flanger mode. This screenshot shows the setting of the deep flanger preset—high feedback, 1 ms of delay, high depth, slow rate, and both the feedback signal and the wet signal are phase-inverted.

Phasing

Flanging and phasing are twin effects, sounding very similar, yet different. A true phasing effect cannot be achieved using a delay. Instead, phasers employ a series of all-pass filters. These filters, instead of altering the amplitude of specific frequencies, change their phase. When the original and treated signals are combined, the result is a series of peaks and dips along the frequency spectrum. But, unlike flanging, the peaks and dips are not organized in harmonic series (Figure 23.4).
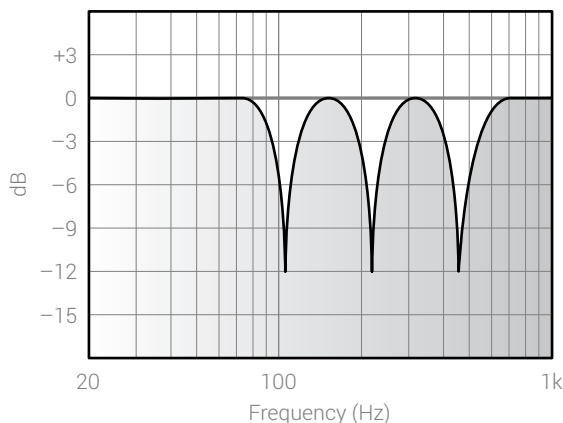
Figure 23.4 The response caused by a sixth-order phaser. The peaks and dips are not harmonically related. As the effect is modulated, this pattern would not expand or contract; it would simply shift lower and higher in frequencies.

Generally, two all-pass filters are required to create one dip. Often phasers use an even number of filters, where fourth order would create two dips, the sixth order would create three, and so forth. The higher the order, the stronger the effect. A feedback loop might also be included, which, like a flanger, would produce a more resonant, bolder effect.

Figure 23.5 The Logic *Phaser* plugin.

Like flanging, phasing is an easily recognizable effect, although generally subtler. In addition to electric guitars, it can be applied on various instruments just to create some cyclic movement or be introduced on specific sections of the song just to add interest.

The settings for all these tracks are identical: fourth order, 0.5 Hz modulation, sweep floor 40 Hz, sweep ceiling 10 kHz. It is interesting to compare these tracks to the flanger tracks and learn to recognize the differences between the two effects.

Track 23.25: Phaser (Vocal)

Track 23.26: Phaser (Distorted Guitar)

Track 23.27: Phaser (Drums)

Track 23.28: Phaser (Shaker)

Plugin: Logic Phaser

Tremolo

A tremolo produces cyclic modulation of amplitude, which makes instruments passing through it periodically fall and rise in level. This effect can be traced back to the Leslie cabinet (mentioned earlier), which also produces a vibrato. While vibrato is a modulation of pitch, tremolo is a modulation of amplitude. The Leslie cabinet (and similar products) is famous for its use with the Hammond organ. Tremolos were later introduced into many guitar amplifiers, and there are endless examples of songs with guitar tremolos. To this day, tremolo is frequently applied on organs and electric guitars, but apart from its legacy there is nothing to prevent it from being applied on any other instrument. Being a distinctive effect in its own right, tremolo can be used either subtly or more evidently to add a sense of motion to sounds.

Figure 23.6 The MOTU *Tremolo* plugin.

Tremolo controls are mostly modulation-related. We often get:

• **Rate**—the modulation rate. Can be expressed in Hz, or if it can be tempo-synced it might be given as note value.
• **Depth**—the intensity of attenuation. Expressed in percentages. Whereas 0 percent means no effect, 100 percent means the signal level will fall down to complete attenuation.

Some tremolos provide a choice of modulation waveform (sine, sawtooth, square). Also, some stereo tremolos employ a separate tremolo for each channel, where the modulation phase between the two can be altered. For example, 180° on a sine modulator would mean that as one channel is at the highest amplitude, the other is at the lowest. This can be used to create interesting panning effects.

Track 23.29: Tremolo 1 (Distorted Guitar)
For the most obvious effect, the tremolo depth in this track is set to 100 percent; the rate is set to 2 Hz.

Track 23.30: Tremolo 1 (Distorted Guitar)
The more musical settings in this track are the outcome of 30 percent depth and 5 Hz rate.

Track 23.31: Tremolo (Vocal)
Forty percent of depth, 5 Hz rate, applied on vocal.

Plugin: UAD Tremolo

24 Reverbs

With the exception of pure orchestral recordings, reverbs are used in nearly every mix. The common practice of close-miking and the dry nature of some sounds produced by synthesizers and samplers result in initial mixes that lack both ambiance and depth (such mixes are often described as "2D mixes"). Reverbs allow us to add and build up these missing elements during mixdown, but they can also be beneficial for many other mixing tasks.

What is reverb?

In nature, reverb is heard mostly within enclosed spaces, such as rooms. Reverbs are easier to understand if we imagine an impulse sound, such as a handclap, emitted from a sound source in an empty room. Such a sound will propagate in a spherical fashion, and for simplicity we should regard it as traveling in all directions. The emitted sound will travel in a direct path to a listener (or a microphone) followed by reflections that bounce from the walls, floor, and ceiling. These will gradually be followed by denser reflections that have bounced many times from many surfaces. As sound diminishes both when traveling through air and being absorbed by surface materials, the reflections will slowly decay in level. Reverb is the collective name given to the sound created by bounced reflections from room boundaries (which we consider to be the main reverb contributors, although in a room there might be many surfaces). In mixing, we use reverb emulators, either hardware units or software plugins, to simulate this natural phenomenon.

Applications

Simulating natural, or creating imaginary, ambiance

Due to the inflexible nature of ambiance recordings and the poor reverb characteristics of many live rooms and home studios, many engineers choose to record most instruments dry and apply artificial ambiance during mixdown. While mixing, reverb emulators give us more options and control over the final ambiance of the mix.

Creating the ambiance of a production is considered one of the more important and challenging practices in mixing—ambiance transforms a collection of dry recordings

into an inspiring spatial arrangement that gives our mix much of its character. The reverb controls let us mold the ambiance and make creative decisions that can make our mix more cohesive and vigorous.

Many regard the creative process of crafting ambiance a vital part of mixing.

Here again, we have a choice between a natural or unnatural (or imaginary) outcome. A natural ambiance is expected in some mixes and for some genres; jazz, bossa nova, and productions that include an orchestral performance, for example. For such projects, we might choose to utilize natural reverb simulations so that the listener can conceive the performance as if being played in a very familiar type of space. Realism is obviously an important factor in natural simulations—not only is the reverb choice important, but also the way we configure it to support the depth of individual instruments and the different properties of the simulated space.

While we can imagine the typical ambiance of a jazz performance, what ambiance should we employ for a trance production? How about electro, hip-hop, trip-hop, techno, or dubstep? There is a loose link between sequenced productions and acoustic spaces. For such projects, the ambiance is realized based on the creativity and vision of the mixing engineer. Furthermore, sometimes an imaginary ambiance is applied even when a natural one is expected. This is done in order to achieve something refreshing, different, and contemporary. Such implementation is evident in some recent jazz albums that, in addition to unnatural ambiance, utilize many modern recording and mixing techniques.

An important point to remember is that while selecting an ambiance reverb, we might ask ourselves, **"How attractive is this reverb?"** and **"How well does it fit into the mix?"** rather than dealing with the question of, **"How natural?"** Such thinking promotes less natural reverbs, and we have already discussed how a non-natural mix can bring about a fantastic result—an imaginary space can excite us more than the familiar, natural ones.

While a natural choice might be suitable for some productions, an imaginary ambiance can sometimes be more rousing and effective.

Gel the instruments in the mix

After combining the various tracks, each instrument in the mix might have a distinctive appearance, but heard together they might feel a little foreign to one another—there might be nothing to link them. A reverb, even one that can hardly be perceived, can gel the various instruments, making them all appear as a natural part of the mix. Similarly to ambiance reverbs, one emulator might have many tracks sent to it, but, in contrast to ambiance reverbs, the gelling reverb is more subliminal: it might not be audible or have a clear spatial sense, and will be felt rather than heard.

Increase the distinction of instruments

Although we have just advocated sending many instruments to the same reverb, different reverbs on different instruments can sometimes be beneficial. For example, having the guitars feeding a specific reverb and the vocals feeding another, despite

creating a less natural space, would increase the distinction between the two instruments, resulting in increased separation.

Depth

In nature, many factors contribute to our perception of depth. The natural reverb in a room is key. During mixdown, reverbs are the main tool we use when positioning sounds in the depth field. This is extremely desirable in many situations—it adds another dimension to the mix, lets us position instruments more naturally, establishes hierarchy, and can also help with masking. But the bond between depth and reverbs does not come without a price—the perceived depth caused by the addition of reverb is sometimes unwanted, for example when adding reverb to vocals while still wanting to keep them in the front of the mix.

> *Although sometimes unwanted, reverb addition often increases the perceived source-to-listener distance, and therefore reverbs are regularly utilized for applying depth.*

Enhance a mood

Different productions have different moods—happy, mellow, dark, and aggressive are just a few examples. Reverbs are often associated with tenderness, romance, mystery, intimacy, and many other sensations and atmospheres. We can enhance, and even establish, a production mood by using reverbs creatively.

Some genres, such as chill-out and ambient, embrace reverbs as a production means. But while reverbs are usually more evident in slow productions, it is perhaps the spatial disorientation caused by certain drugs that fosters reverb use in some upbeat dance genres (and psychedelic music in general).

Livening sounds

Those who have heard a recording that took place in an anechoic chamber or those who have attended an unamplified performance in an open field know how unsatisfying such a reverb-free performance can be. Reverb is a native part of our life as it accompanies most of the sounds we are exposed to. A reverb-free recording of specific instruments, as good as it may be, can sound cold and unnatural. By adding reverb, we take a dry, artificial recording and breathe "life" into it—adding the reverb we are so familiar with in daily life.

However, it is not only the existence of a reverb that matters—it is also the amount of reverb. Ask yourself whether you like your singing better when "performing" in your bedroom or in your bathroom. Most people will choose the latter since mostly it has a louder and longer reverb. Reverbs can make sounds more impressive and more present; when we say "livening," we also mean taking a stodgy performance and turning it into a captivating one.

Another issue that comes to mind here is the use of mono samples in sequenced production. For instance, a mono ride sample can sound odd compared with the

familiar sound of a ride captured by stereo overheads. Adding reverb to mono samples can produce a more natural and less mechanical feel to sequenced instruments.

Filling time gaps

Empty time gaps can break the balance and integrity of our mix. Reverbs can make mix elements longer and fill empty time gaps, an application more common in slow, sparse productions.

Filling the stereo panorama

Consider a slow production consisting of an acoustic guitar and vocal recordings only. How can one craft a full stereo panorama with just two instruments? Stereo reverbs can be utilized in mixes to create an imaginary sound curtain behind the performers. Often reverbs of such a type span the stereo panorama. But we can also use reverbs to fill gaps in the stereo panorama by simply padding specific instruments with a stereo reverb, making their stereo image wider.

Track 24.1: Guitar Dry
The source guitar track, dry, with its focused and thin image at the center of the panorama.

Track 24.2: Guitar with Reverb
The added reverb fills the stereo panorama. Also, notice that the guitar appears farther back, although the level of the dry signal has not changed.

Plugin: Universal Audio DreamVerb

Changing the timbre of instruments

In many situations, added reverbs modify the timbre of the original sound, whether it be our desire or not. A loud reverb on a percussive instrument can smoothen its attack, slow its decay, or create a completely different decay envelope. The friction between a violin string and a bow can produce a harsh noise that can be blurred and softened by a reverb. Conversely, though, it is exactly that friction noise that can give a double bass added expression in the mix. A reverb can be used to constructively alter timbre, but it can also deform it (a possibility we must be aware of). In most situations, the right reverb choice along with appropriate control adjustment will render the desired result.

> *When adding reverb, we listen closely to timbre changes and, when*
> *needed, tweak the reverb parameters.*

Reconstruct decays and natural ambiance

Reverbs can also be used to replace the natural decay of an instrument. Such an application might be required if a tom drum recording contains exaggerated spill. Toms are commonly compressed in order to give them more punch and bolder decay, but

doing so will also amplify the spill, which can cause some level fluctuations in the mix. To solve this, we gate the toms so the decay is shortened, and then feed the gated tom into a reverb unit, which produces an artificial decay. We have much more control over artificial decays and we can, for example, make them longer and bigger, so as to make the toms prominent.

A similar application is the reconstruction of room ambiance. A classic example is the ambiance captured by a pair of overheads. It can be very tricky to control such ambiance and fine-tune it into the mix, mainly with regards to frequencies and size. It can become even more problematic after these overheads tracks are compressed, since this will emphasize the ambiance even more. The low and mid-frequencies, where most of the presence and size of the ambiance are found, can be filtered and reconstructed using a reverb.

Track 24.3: Drums Before Reconstruction
These are the original drums, unprocessed.

Track 24.4: Drums After Filtering
Only the room and the overheads are filtered in this track. A 12 dB/oct HPF at 270 Hz removes much of the ambiance. The drums now sound thinner and their image appears to be closer.

Track 24.5: Drums After Reconstruction
The close-mics are sent to a reverb emulator, which is preceded by an LPF. The result is ambiance different in character from that in Track 24.3.

Plugin: Audio Ease Altiverb

Drums: Toontrack EZdrummer

Resolve masking

Although a potential masker in its own right, a reverb can be used to resolve masking. The common usage of reverbs for the purpose of depth gives us the ability to position sounds in a front–back fashion, in addition to the conventional left–right panning. This adds another dimension and much more space in which we can place our instruments.

A stereo reverb, which spreads across the sound stage, has more chance of escaping masking, and its length means that it will have more time to do so. Say we have two instruments playing the same note and one completely masks the other. By adding reverb to the masked instrument, we make it longer and are able to identify it more because of its reverb.

More realistic stereo localization

Sounds that are not accompanied by reflections are very rare in nature. Panning a mono source to a discrete location in the stereo panorama can therefore produce an artificial

result. Reverbs can be added to mono signals in order to produce more convincing stereo localization. The reverb will also extend the stereo width of the mono signal, which can make it clearer and more defined. The reverbs that are used for such tasks must be stereophonic and rely on a faithful early reflection pattern, and are usually very short.

A distinctive effect

Sometimes reverbs are added just to spice things up and add some interest. Snare reverb, reversed reverb, gated reverb, and many other reverbs are simply added based on our creative vision and not necessarily based on a practical mixing need. Reverbs can be automated between sections of the production in order to create some movement or be introduced occasionally to mark transitions between different sections.

Impair a mix

Although not exactly a constructive application, it is worth looking at the common problems that reverbs might cause and what we should watch out for:

- **Definition**—reverbs tend to smear sounds and to make them unfocused and distant, and can decrease both intelligibility and localization. Proper configuration of a reverb emulator can usually prevent these problems.
- **Masking**—as they are usually long, dense, and wide sounds, reverbs can mask other important sounds. As with any case of masking, mute comparison can reveal what and how reverbs are masking.
- **Clutter**—if used unwisely (which usually means too loud and for too long), reverbs can clutter the mix and give a "spongy" impression. Part of the challenge is making reverbs effective without being dominant.
- **Timing**—as a time-based effect, a reverb can affect the timing of a performance, especially a percussive one.
- **Change timbre**—discussed previously.

Types

Stereo pair in a room

The early method of incorporating reverbs into a mix (long before the invention of reverb emulators) was by capturing the natural reverb using a set of microphones in a room. Producers used to position musicians in a room in relation to how far away and how loud they should appear in the mix. Back then, and even today, different studios are chosen based on their natural reverb characteristics. In some studios, special installations, such as moving acoustic panels, are used to enable a certain degree of control over the reverb.

In most situations where a small sound source is recorded, say fewer than ten musicians, the ambiance and depth of the room are captured by two microphones positioned using an established stereo-miking technique. As already discussed, such a stereo recording results in what can be described as a submix between the instrument

and its reverb, which can be altered only slightly during mixdown. However, the celebrated advantage of this method is that it captures the highest extent of complexity that a natural reverb offers. No reverb emulator has the processing means to produce such a faithful, and natural, simulation, especially not in real time. As acoustic instruments radiate different frequencies in different directions, a microphone in close proximity can only capture parts of the timbre. But all radiated frequencies reflect and combine into a rich and faithful reverb, which is then captured by the stereo pair.

The reverb and depth captured by a pair of microphones in a room are the most natural and accurate ones, but they can be seriously limiting if they do not fit into the mix.

Reverb chambers (echo chambers)

A reverb chamber is an enclosed space in which a number of microphones are placed in order to capture the sound emitted from a number of speakers and the reverb caused by surface reflections. Either purpose-built or improvised, these spaces can be of any size and shape. Even a staircase has been utilized for such a purpose. A send from the control room feeds the speakers in the reverb chamber and the microphone signals serve as an effect return. While reverb chambers are generally not an accessible mixing tool for most engineers, every studio can be used as a reverb chamber—all studios (live rooms) have microphone lines and many have loudspeakers installed with a permanent feed from the control room.

The advantage of using reverb chambers is that the captured reverb is, again, of a very high quality, being produced by a real room. As opposed to a stereo recording, in which most mixing elements are immovable, reverb chambers are used during mixdown and provide more control over the final sound. This is achieved either by different amounts of send and return or by physical modifications applied in the chamber itself; for example, changing the distance between the microphones and the speakers or adding absorbent (or reflective) materials. However, such alterations result in a limited degree of control, and most of the reverb characteristics will only vary by insignificant amounts, most notably the perceived room size. We consider each reverb chamber as having a very distinctive sound that can only be "flavored" by alterations. Moreover, the loudspeakers used in a reverb chamber differ substantially from a real instrument in the way they produce and radiate sound, so reverb chambers are just a simple alternative to placing a real instrument in a real room. Another issue with reverb chambers is their size—a general rule in acoustics is that larger rooms have better acoustic characteristics than smaller ones. A small reverb chamber can produce a reverb with broken frequency response and reflections that can color the sound. A large reverb chamber is expensive to build and only big studios can afford it.

Track 24.6: Vocal Chamber Convolution
A chamber reverb simulation produced by a convolution emulator.

Plugin: Trillium Lane Labs TL Space

Spring reverb

Since reverb is essentially a collection of delayed sound clones, it is no wonder that one of the first artificial reverbs to be invented had its origins in a spring-based delay device. The original device was conceived by Bell Labs researchers who tried to simulate the delays occurring over long telephone lines. The development of the spring reverb, starting as early as 1939, is credited to engineers from Hammond who tried to put some life into the dry sound of the organ. During the early 1960s, Leo Fender added Hammond's spring reverb to his guitar combo and was followed by manufacturers such as Marshall and Peavey.

A spring reverb is an electromechanical device that uses a system of transducers and steel springs to simulate reflections. The principle of operation is simple: an input transducer vibrates with respect to the input signal. Attached to the transducer is a coiled spring on which the vibrations are transmitted. When vibrations hit the output transducer on the other end of the spring, they are converted into output signals. In addition, parts of the vibrations bounce back onto the spring and then bounce back and forth between the ends. Such reflections are identical to those bouncing between "surfaces in a one-dimensional space" (technically speaking, a line, with its boundaries being its end points).

In reality, the science of a spring reverb is more complex than the above explanation, but these units are relatively cheap to manufacture and still ship today with many guitar amplifiers. Although spring reverbs do a pretty bad job in simulating a natural reverb, listeners have grown to like their sound. Furthermore, many digital reverb emulators fail to match the sound of a true electromechanical spring reverb, which explains why standalone rack-units are still in production. Such units will be fitted with very few controls and in most cases will have fixed decay and pre-delay times. (It is worth noting that pre-delay could always be achieved if a delay was connected before the reverb; however, it was not until the early 1970s that digital delays started replacing tape delays for such purposes.) Quiet operation and flat frequency response were never assets of true electromechanical spring reverbs, and they are known to produce an unpleasant sound when fed with transients. Therefore, spring reverbs are usually applied on the likes of pads and vocals, normally while only mixing a small amount of the wet signal.

Figure 24.1 The Accutronics *Spring Reverb*.
Source: Courtesy of Sound Enhancement Products, Inc.

While spring reverbs are far from spectacular in simulating a natural reverb, the sound created by a true electromechanical unit can serve as an identifiable retro effect, especially if applied on guitars and other non-percussive instruments and sounds.

Track 24.7: Vocal Spring Convolution
A spring reverb simulation produced by a convolution emulator.

Track 24.8: Snare Spring Convolution
The snare on this track is sent to the same emulator and preset as in the previous track.

Plugin: Trillium Lane Labs TL Space

Plate reverb

If a spring reverb is considered a "one-dimensional" reverb simulator, the logical enhancement to such a primitive design was going "two-dimensional." The operation of a plate reverb is very similar to that of a spring reverb, only that the actual vibrations are transmitted over a thin metal plate suspended in a wooden box. An input transducer excites the plate and output transducers are used to pick up the vibrations. In the case of plates, these vibrations propagate on its two-dimensional surface and bounce from its edges.

In 1957, a pioneering German company called EMT invented and built the first plate reverb—the *EMT 140*. This model, which is still a firm favorite, had the impressive dimensions of 8 ft × 4 ft (2.4 m × 1.3 m) and weighed around 600 pounds (more than a quarter of a ton). These original units were very expensive but still much cheaper than building a reverb chamber. They can be heard on countless records from the 1960s and 1970s and on other productions that try to reproduce the sound of these decades.

While mobility was not one of its main features, the plate reverb had better sonic qualities than the spring reverb. In addition to a damping mechanism that enabled control over the decay time, its frequency response was more musical. The reverb it produced, although still not resembling a true natural reverb and being slightly metallic, blended well with virtually every instrument, especially with vocals. The bright, dense, and smooth character of plate reverbs also made them a likely choice for drums, which explains why they are so frequently added to snares.

Although a serious artificial reverb simulator, a plate reverb does not produce a truly natural-sounding reverb. Hence, it is used more as an effect to complement instruments such as snares and vocals.

Track 24.9: Vocal Plate Convolution
A plate reverb simulation of the EMT 140 produced by a convolution emulator.

Track 24.10: Snare Plate Convolution
The same emulator and preset as in the previous track, only this time being fed with a snare.

Plugin: Trillium Lane Labs TL Space

Digital emulators

So far, we have discussed natural and mechanical reverbs. The invention of digital reverberation is credited to Manfred Schroeder, then a researcher at Bell Laboratories, who, back in 1961, demonstrated a simple digital reverberation system (Figure 24.2a). It took a while to turn his digital reverberation ideas into a tangible commercial machine, but with the performance rise of DSP chips and the fall in their price, digital emulators were destined to take over. It was EMT again, with help from an American company called Dynatron, that, in 1976, revealed the world's first commercial reverb unit—the *EMT 250*. It consisted of very basic controls, such as pre-delay and decay time, and produced effects such as delay and chorus.

A digital reverb is not a dense collection of simple delays, since these would cause much coloration to the direct sound and produce a broken frequency response. There are many internal designs for digital reverbs and they vary in their building blocks and in the way these blocks are connected. At the lowest level, digital emulators are implemented using mathematical functions that are executed on internal DSP chips. To differentiate from digital convolution reverbs (discussed later), these types of reverbs are now referred to as *algorithmic reverbs*.
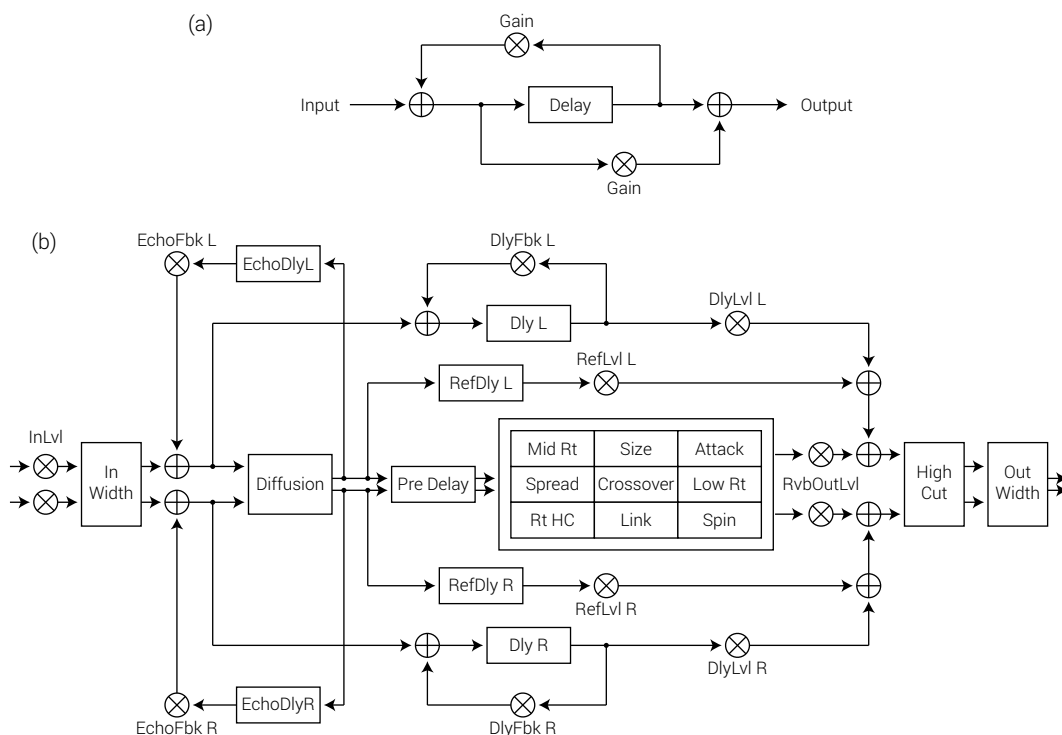
Figure 24.2 Block diagram of a reverb. (a) The original Schroeder reverb design (1961). Later digital designs are far more complex, as can be seen in (b): the plate program of the Lexicon *PCM91* (1998).

Track 24.11: VocalEMT 250 Convolution
A plate reverb simulation of the EMT 250 produced by a convolution emulator. Notice that, compared to Track 24.9, this reverb has better spatial characteristics.

Plugin: Trillium Lane Labs TL Space

No digital emulator will ever be able to produce a reverb completely identical to the one created in a real room. This is mainly due to the complexity of such a reverb—there are thousands of reflections to account for; different frequencies propagate and diffract in different fashions; different materials diffuse and absorb sound in different ways; and even the air itself changes the sound as it travels through it. Manufacturers must take shortcuts, and the more processing power that is at their disposal the fewer shortcuts they have to take, and thus the more realistic the reverb is likely to be. It is worth noting that although DSP chips were far less powerful in the 1970s and 1980s, many units from these decades are still regarded today as state of the art, which shows the importance of the design of such units. Since each respectable manufacturer has their own secret architecture, each has a distinctive sound.

Back in the 1990s, when real-time plugins emerged, CPUs had less than a tenth of the processing power compared with modern processors and could only handle a few plugins at a time. Today, more and more people rely solely on their computer CPU for all mixing tasks. Yet, a high-quality reverb plugin can consume significant amounts of the CPU processing capacity, so many still find the need for additional
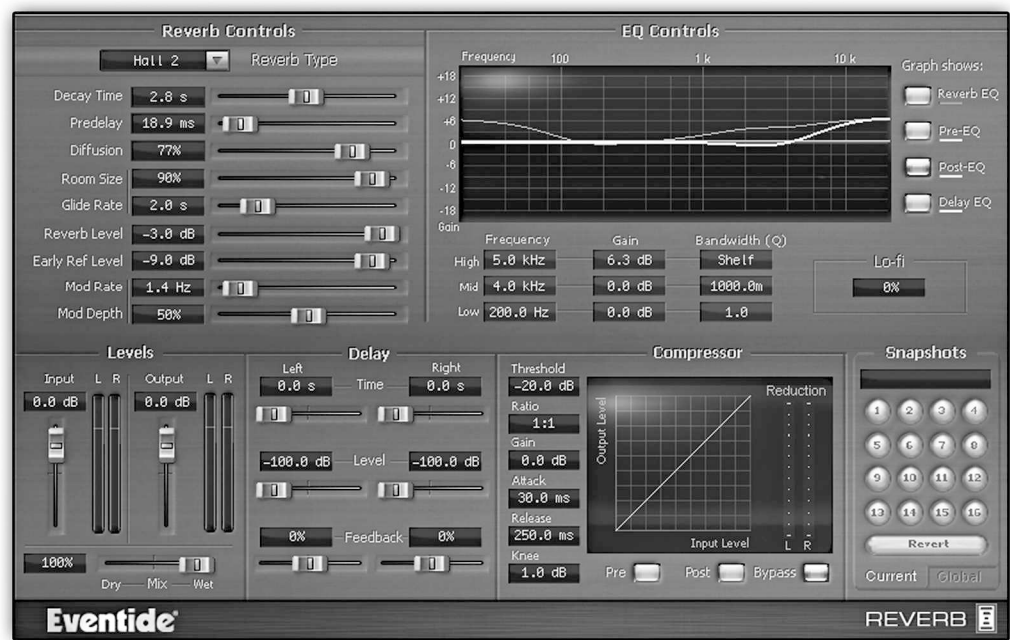
Figure 24.3 An algorithmic reverb plugin—the Eventide *Reverb TDM*.

Figure 24.4 The "larc" (lexicon alphanumeric remote control) for the Lexicon *480L*; the reverb unit itself is rack-mounted and often unseen in studios. Despite being released in 1986, the *480L* is still considered by many as one of the best reverb emulators.
Source: Courtesy of SAE, London.

processing power—either in the form of external hardware units or as internal DSP expansion cards.

As algorithmic reverbs have no physical or mechanical limitations, they provide a multitude of controls that let us tweak nearly every property of the reverb. This makes them an extremely flexible and versatile tool in many mixing scenarios, and it is no wonder that they are the most common reverbs in our mixing arsenal. An issue with high-quality emulators is that they are expensive both in terms of price and processing-power consumption.

> *Digital reverbs are the most common type in use and give us great control over reverb parameters, but they can consume large amounts of processing power, depending on their quality.*

Convolution (sampling) reverbs

Even back in the 1970s, various people toyed with the possibility of capturing the reverb characteristics of an acoustic space, so it could be applied later to any kind of recording. What may have sounded like a pipe dream has become a reality with the new millennium, with DSP technology now fast enough to accommodate the real-time number of calculations required for such a process.

Reverb sampling is normally done by placing a pair of microphones in a room, and then recording the room response (i.e., the reverb) to a very short impulse, like that of a starter pistol. Since it is very hard to generate a perfect impulse, an alternative

method involves playing a sine sweep through speakers instead. The original sound might be removed from the recording, leaving the reverb only. The recorded impulse response (IR) is then loaded into a convolution reverb, which analyzes it and constructs a huge mathematical matrix that is later used by the convolution formula. With every sound fed through the unit, a reverb very similar to that of the original space is produced.
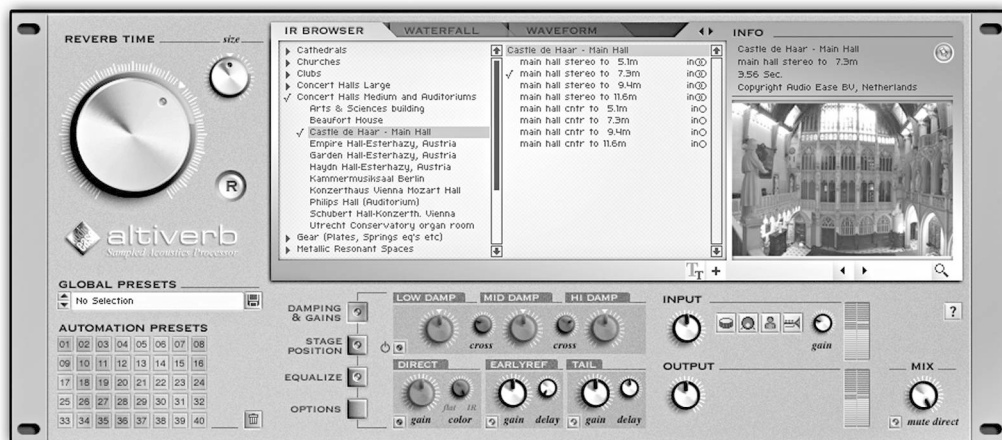
Figure 24.5 A convolution reverb plugin, the Audio Ease *Altiverb*.

An emulator can be based on one of two types of convolution: either one that is done in the time domain (pure convolution) or one that is based on the frequency domain (convolution or Fourier-based). Each generates the same result, only in some situations one will be faster than the other. If pure convolution were used, an impulse response of 6 seconds at 44.1 kHz would require around 23 billion mathematical operations per second—equivalent to the processing power offered by a 2.2 GHz processor. It can easily be seen how such a process might be unwieldy in some situations. As a general rule, the shorter the original impulse response is, the less the processing needed would be.

Convolution reverbs experience increased popularity nowadays. They let us incorporate into our mix the reverb of many exotic venues and spaces. For the film industry, this is a truly revolutionary tool—engineers can record the reverb of any location and then apply it during post-production. For mixing, however, it is doubtful how much the reverb characteristics of the Taj Mahal can contribute to a modern rock production. Yet, the impulse recordings shipped with emulators include less exotic spaces that can be used in every mix. Many impulse responses can also be downloaded from the Internet; some are free. The quality of the impulse recording is determined by the quality of the equipment used, which is a vital factor if natural results are sought. It is generally agreed that a good impulse recording produces an extremely believable reverb simulation that matches (if not exceeds) the quality of the best algorithmic emulators.

Convolution reverbs are not only used to reproduce the reverb of acoustic spaces. Impulse responses of virtually any vintage unit, such as the *EMT 140* or *250*, are also

available. Fascinating results can also be achieved if instead of loading a real impulse response one loads just a normal short recording.

One problem with convolution reverbs is that all the properties of the reverb, as captured during the recording process, are imprinted into the impulse recording and can be altered very little later on without some quality penalty. An example of this is the inability to change the distance between the sound source and the listener or the different settings used while a hardware unit was sampled. Many convolution reverbs include only limited numbers of controls that make use of simple envelopes and filters in order to give some degree of control over the reverb characteristics. But for more natural results, many reverbs ship with a variety of impulse recordings of the same space, each based on a different recording setup. Generally speaking, convolution reverbs produce the best simulations provided the original impulse response is unaltered by the artificial envelope and filter controls. This fact makes convolution reverbs, to some extent, a hit-and-miss affair—if the pure impulse response is right, it should sound great; if it isn't, additional tweaking might make it more suitable for the mix, but not as great. As with algorithmic reverb, the initial preset choice is crucial.

Good convolution reverbs can produce exceptionally natural results, but they consume processing power and are normally inflexible.

Perhaps the irony of convolution reverbs is that now, after having them at our disposal, many find the old algorithmic reverbs far more interesting and suitable for both creative and practical tasks. Despite the fact that convolution reverbs are considered superior in simulating natural spaces, algorithmic reverbs are still favored for many other mixing applications.

The following tracks demonstrate various IR presets from the Logic Space Designer convolution reverb. It is for readers to judge how realistic and appealing each of these reverbs is. It is perhaps most interesting to observe how, despite hardly adding a noticeable reverb, the Wine Cellar track has an extra sense of space and depth compared to the dry track.

Track 24.12: Conv Source
The dry source track used in the following samples.

Track 24.13: Conv Wine Cellar

Track 24.14: Conv Wooden Room

Track 24.15: Conv Drum Chamber Less

Track 24.16: Conv Wet Bathroom

Track 24.17: Conv Small Cave

Track 24.18: Conv Music Club

Track 24.19: Conv Concert Hall

Track 24.20: Conv Canyon

Track 24.21: Conv Roman Cathedral

Here are two sets of samples demonstrating the use of just a normal sample as an IR. The results are unnatural, but can be used in a creative context:

Track 24.22: Conga IR

Track 24.23: Conv Conga

Track 24.24: FM Noise IR

Track 24.25: Conv FM Noise

Plugin: Logic Space Designer

Reverb programs

Digital emulators can simulate a number of spaces; each will exist as a loadable program (preset). As there can be hundreds of programs, these are organized in categories.

A high-quality reverb emulator might implement a different algorithm for each category. The main categories are:

Category	Description	Application
Halls	Large, natural-sounding, live spaces	Natural
Chambers	Simulate reverb chambers or spaces that have slightly less natural reverb behavior and a less defined size	Natural
Rooms	Normal rooms of different sizes	Natural
Ambiance	Concerned more with placing the sound naturally in a virtual space, caring less about the actual reverberation; most often an ambiance preset involves early reflections only	Natural
Plate	Plate reverbs	Effect

The following categories are also common:

Category	Description	Application
Studio	Simulate the reverb in a recording studio live room	Natural
Church/ cathedral	These types of spaces might produce a highly impressive reverb for certain types of instruments, such as organs, but they generally result in poor intelligibility	Natural
Spring	Spring reverb	Effect
Gated	Nonlinear reverb (explained later)	Effect
Reversed	Rising reverb instead of decaying one (explained later)	Effect

Some mixes are expected to present a more natural space than others. If seeking a truly natural space in the mix, hall, or room, presets can produce good results. If less-natural reverb is required, either chamber or ambiance might be used. Most ambiance programs are essentially the shortest and most transparent of all reverb programs. They excel at blending mono instruments into a spatial arrangement, without adding a noticeable reverb.

The choice of category is usually determined by two main factors. First, the size of the room should fit with the type of music. A classical recording or a chill-out production might use a large space such as a hall; a bossa nova or jazz production can benefit from a moderate-sized hall; while a heavy metal or trance production might make use of very small space simulations. Second, the decay time (or length) of the reverb should be relative to the mix density—a dense mix will suffer from long reverb tails that will cause masking and possibly clutter; in a sparse mix, longer decays can fill empty spaces.

For a truly natural space, hall and room programs can work better than chambers and ambiance. The room size and the decay time are the two main factors when selecting a reverb program.

Within each category, various programs reside. For example, the halls category might include small hall, large hall, vocal hall, and drum hall. As opposed to EQs and compressors, it is uncommon for a mixing engineer to program a reverb from scratch—in most cases, a preset selection is made and then altered if needed. Initially, choosing the right program for the mix will involve experimentation—many mixing engineers will focus on one or two categories and will try all the different programs within each; if different units are available, each of these might be tested as well. It is time well spent, especially when selecting one of the main reverbs in the mix. A reverb that sounds good when soloed might not interact well with the rest of the mix—final reverb selection is better done in mix context. With experiment comes experience, and after a while we learn which emulator and which program works best for a specific task. The factory settings of each program are optimized by the manufacturer based on studies and experts' opinions; therefore, drastic changes to these settings can theoretically degrade the quality of the simulation. This makes the initial selection of an appropriate program even more important.

The selection of an appropriate reverb program is highly important and is a process worth spending time on.

When a reverb is used as an effect, an artistic judgment is required since there are no golden rules. Many use plate programs on vocals and drums; snares are commonly treated with a plate reverb or a gated one. For a more impressive effect, chambers or halls can work. Halls can also be suitable for orchestral instruments such as strings, flutes, brass, and the saxophone. Bass guitars are usually kept dry, and distorted electric guitars can benefit from a subtle amount of small-room reverb that will add a touch of shine and space. As the association between synthesized sounds and natural acoustic spaces is loose, chamber and ambiance programs might be more suitable in a sequenced production.

The above should merely serve as a guideline—the final choice of reverb program should be a result of experimentation and will vary for each individual mix.

The following tracks demonstrate various reverb programs. Reverb programs can vary noticeably between one emulator and another, so readers are advised to treat the following tracks as a rough sample.

Track 24.26: Vocal Dry

Track 24.27: Drums Dry
The source tracks used in the following samples.

Track 24.28: Program Hall (Vocal)

Track 24.29: Program Hall (Drums)

Track 24.30: Program Chamber (Vocal)

Track 24.31: Program Chamber (Drums)

Track 24.32: Program Room (Vocal)

Track 24.33: Program Room (Drums)

Track 24.34: Program Ambiance (Vocal)

Track 24.35: Program Ambiance (Drums)

Track 24.36: Program Plate (Vocal)

Track 24.37: Program Plate (Drums)

Track 24.38: Program Studio (Vocal)

Track 24.39: Program Studio (Drums)

Track 24.40: Program Church (Vocal)

Track 24.41: Program Church (Drums)

Track 24.42: Program Cathedral (Vocal)

Track 24.43: Program Cathedral (Drums)

Track 24.44: Program Spring (Vocal)

Track 24.45: Program Spring (Drums)

Plugin: Audio Ease Altiverb

Reverb properties and parameters

The parameters found on most digital reverb emulators are closely related to the properties of the natural reverb produced in acoustic spaces. The understanding of these parameters is important not only for more-natural reverb simulations, but also for incorporating reverbs more effectively and musically into mixes.

Reverb emulators vary in their internal design, and some are intended for specific applications more than others. The controls offered by reverb emulators can vary substantially between one emulator and another, and it is impractical to discuss all the possible controls offered by each individual emulator. Moreover, identical controls on two different emulators can be implemented in different ways and can yield different results. Reading the manual will reveal the purpose of each control, and some manuals include useful tips on optimal utilization of the emulator or reverbs in general. This section covers the most important parameters that are typical of many emulators.
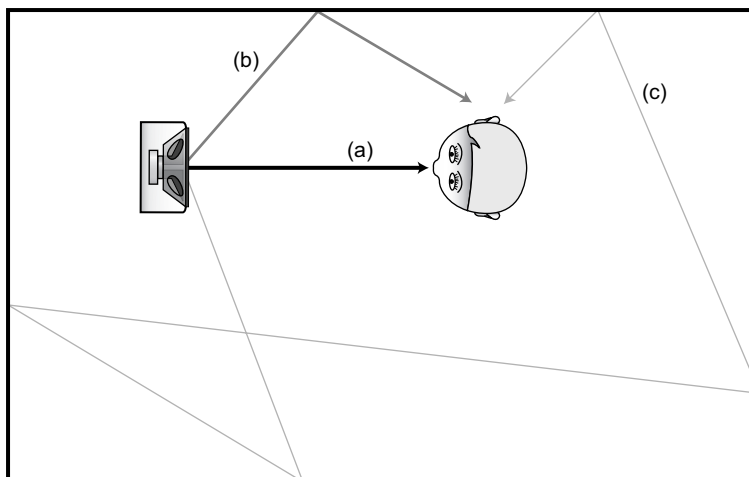
Figure 24.6 Direct and reflected sounds. (a) The direct sound reaches our ears without encountering any boundaries. (b, c) Reflected sound. Note that the path of reflection (c) is valid, considering the propagation characteristics of low frequencies.

Direct sound

Direct sound is not part of a reverb. It is the sound that travels the shortest distance between the sound source and the listener, and in most cases it does so while traveling in a direct line between the two [path (a) in Figure 24.6]. It is the first instance of the sound arriving at the listener and hence provides an important psychoacoustic cue. As discussed in previous chapters, both the level and the high-frequency content of the direct sound contribute to our perception of depth.

The direct sound is the dry signal that we feed into a reverb emulator so it can produce a simulation of a reverb. A reverb emulator might have a parameter to determine whether this dry signal is mixed with the generated reverb (dry/wet mix). If the reverb is connected via an aux send, the original dry signal is mixed anyway, and the copy sent to the emulator becomes redundant once mixed with the reverb. This is not the case when a reverb is connected as an insert, where the dry signal is only heard if blended with the reverb at the emulator output.

If reverbs are connected via an aux send, make sure to toggle off the dry signal on the emulator. However, toggle it on if reverbs are connected as an insert.

Pre-delay

Pre-delay is the time difference between the arrival of the direct sound and that of the very first reflection. However, in a few reverb emulators, this parameter stands for the time gap between the dry signal and the later reverberation.

Pre-delay gives us a certain clue to the **size** of the room; in larger rooms, the pre-delay is longer as it takes more time for reflections to travel to the boundaries and back to the
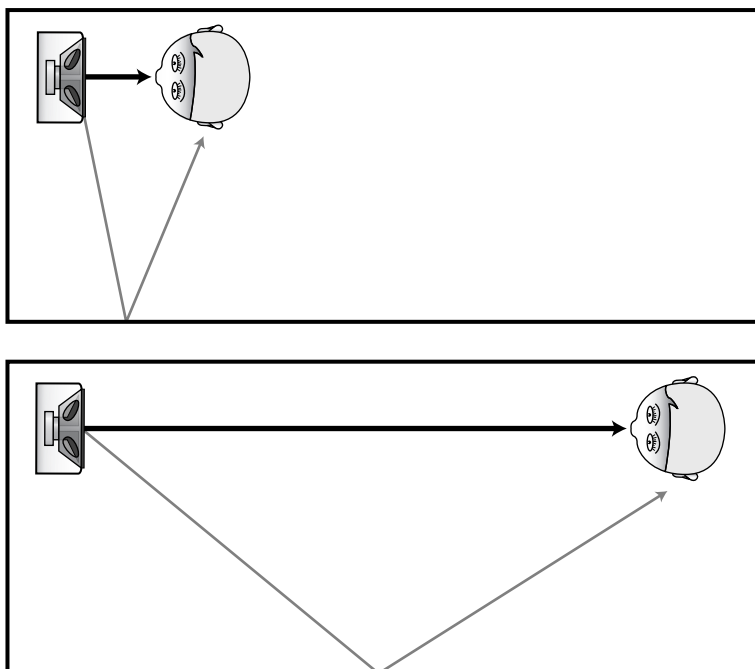
Figure 24.7 Distance and pre-delay. When the listener is closer to the sound source, the relative travel distance between the direct sound and the first reflection is bigger, which results in a longer pre-delay.

listener. Pre-delay also gives a clue to the **distance** between the source and the listener, but here the opposite of what we might initially expect takes place: the closer the source to the listener, the longer the pre-delay. This is due to the fact that the *relative* distance between the direct and the reflected sounds is getting smaller the farther away the source is from the listener (Figure 24.7). This phenomenon is commonly put into practice when we require a reverb but not the depth that comes with it—we simply lengthen the pre-delay. Pre-delay is normally expressed in milliseconds, and for natural results our brain requires that it is kept below 50 ms. However, longer pre-delay times are still used sometimes.

> *A long pre-delay time can help keep an instrument at the front of the mix.*

The pre-delay time also determines when reflections start to mix with the direct sound. Reflections caused by a real room are far more complex than those produced by an emulator with its limited building blocks. Needless to say, high-quality emulators do a better job in that respect, but we still say that digital reflections are more likely to color the original sound than those created by a real room. Thus, the reflections generated by a digital emulator might cause comb filtering and other side effects when mixed with the original signal. The sooner they are mixed, the more profound this effect will be. It is worth remembering that even the sound of a snare hit can easily last for 80 ms before entering its decay stage. If the early reflections are mixed with the direct sound within the very first milliseconds (like they mostly do in nature), the timbre might distort. We lengthen the pre-delay time in

order to minimize this effect. On the same basis, early reflections can be masked by the original signal. Since the early reflections give us most of the information regarding the properties of the space, lengthening the pre-delay can nudge these reflections outside the masking zone and result in clearer perception of important psychoacoustic cues.

Another issue related to short pre-delay settings is intelligibility—if the reverb is mixed with direct sound very early, it might blur and harm both clarity and definition. This is extremely relevant when adding reverb to vocals, as in most cases the lyrics need to be intelligible.

A very long pre-delay can cause an audible gap between the original and the reverb sounds. This will separate the reverb from the dry sound, cause the reverb to appear behind the dry material, and usually produce an unnatural effect along with possible rhythmical disorder. We usually aim at a pre-delay time that will produce minimal timbre distortion without audible time gaps. Nevertheless, audible pre-delay gaps have been used before in order to achieve interesting effects, including rhythmical ones.

> *A too-short pre-delay can distort timbre, while too long can*
> *cause audible time gaps. We usually aim for settings that*
> *exhibit none of these issues while taking into account*
> *the intelligibility of the source material.*

Track 24.46: Pre-Delay 0 ms Snare
When the pre-delay is set to 0 ms, both the dry snare and its reverb appear as one sound.

Track 24.47: Pre-Delay 50 ms Snare
The actual delay between the dry snare and the reverb can be heard here, although the two still appear to be connected.

Track 24.48: Pre-Delay 100 ms Snare
One hundred milliseconds of pre-delay cause two noticeable impulse beats. The obvious separation between the two sounds yields a non-musical result.

Track 24.49: Pre-Delay 0 ms Vocal
As with the snare sample, a 0 ms delay unifies the dry voice and the reverb into a smooth sound.

Track 24.50: Pre-Delay 25 ms Vocal
Despite a hardly noticeable delay, the intelligibility of the vocal is slightly increased here. This is due to the fact that the delayed reverb density leaves some extra space for the leading edge of the voice.

Track 24.51: Pre-Delay 50 ms Vocal
The 50 ms delay can be discerned in this track, and it adds a hint of a slap-delay. Similarly to the previous track, the voice is slightly clearer here as there is more space for the leading edge of the voice. However, the delay makes the overall impression busier and less smooth, resulting in slightly more clutter.

Track 24.52: Pre-Delay 100 ms Vocal
One hundred milliseconds of pre-delay clearly results in a combination of a delay and a reverb. This is a distinctive effect in its own right, which might work in the right context. Notice how, compared to the 0 ms version, the reverb on this track appears clearer and louder.

Plugin: t.c. electronic ClassicVerb

Early reflections (ERs)

Shortly after the direct sound, bounced reflections start arriving at the listener [path (b) in Figure 24.6]. Most of the early reflections only bounce from one or two surfaces, and they arrive at relatively long intervals (a few milliseconds). Therefore, our brain identifies them as discrete sounds correlated to the original signal. The early reflections are indispensable in providing our brain with information regarding the space characteristics and the distance between the source and the listener. A faithful early reflection pattern is obligatory if a natural simulation is required, and alterations to this parameter can greatly enhance or damage the realism of the reverb. Neither spring nor plate reverbs have distinct early reflections due to their small size, and therefore they do not excel at delivering spatial information.

Gradually, more and more reflections arrive at the listener, and their density increases until they cannot be distinguished as discrete echoes. With dependence on different room properties, early reflections might arrive within the first 100 ms following the direct sound. It is worth remembering that early reflections within the first 35 ms fall into the Haas zone; hence, our brain discerns them in a slightly different manner. In addition, these very early reflections are readily masked by the direct sound. As discussed earlier, we can increase the pre-delay time in order to make the early reflections clearer.

Early reflections provide our brain with most of the information regarding space properties and will contribute greatly to the realism of depth.

The **level** of the early reflections suggests how big the room is—a bigger room will have its boundaries farther away from the listener, and thus bounced reflections will travel longer distances and will be quieter. **Surface materials** also affect the level of reflections. For example, reflections from a concrete floor would be louder than reflections from a carpeted floor.

In relation to **depth**, the level of early reflections might again have the opposite effect to what we might expect. Although the farther away the listener is from the sound source the longer distance the reflected sound travels, it is the difference in distances traveled between the direct and the reflected sounds that matters here—a close sound source will have a very short direct path but a long reflected path. The farther away the source is from the listener, the smaller the difference in distance between the two paths. In practice, the farther away the source and the listener are, the closer in level the direct and reflected sounds will be; or, in relative phrasing, the louder the early reflections will be.
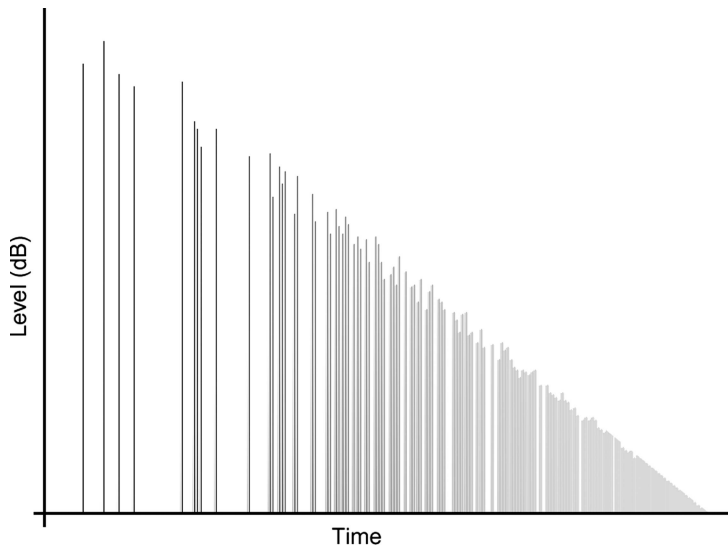
Figure 24.8 Possible reflection pattern in a real room.

Louder early reflections denote greater distance between the source and the listener.

Early reflections are the closest sound to the dry sound; hence, they are the main cause for timbre distortion and comb filtering. One of the biggest challenges in designing a reverb emulator involves the production of early reflections that do not color the dry sound. Sometimes attenuating or removing the early reflections altogether can yield better results and a more healthy timbre. The same practice can also work with a multi-performance such as a choir—in a real church, each singer produces a different early reflection pattern, and the resultant early reflection pattern is often dense and indistinct. Nevertheless, removing the early reflections can result in poor localization and vague depth, so it may be unsuitable when trying to achieve a natural sound.

Finally, there is a trick that can be employed to add the early reflections alone to a dry signal. This can enliven dry recordings in a very transparent way and with fewer side effects. The explanation for this phenomenon goes back to the use of delays to open up sounds and create some spaciousness. Good early reflections engines create more complex delays that will not color the sound as much while giving a greater sense of space. This trick can work on any type of material, but is particularly effective on distorted guitars.

Early reflections can distort the dry signal timbre. For less natural applications, concealing the early reflections can sometimes be beneficial. Conversely, sometimes the early reflections alone can work magic.

Track 24.53: Organ Dry
The dry signal used in the following tracks.

Track 24.54: Organ Dry and ER
The dry signal and the early reflections with no reverberation. Note the addition of a spatial cue and the wider, more distant image.

Track 24.55: Organ ER Only
The early reflections in isolation without the dry signal or reverberation.

Track 24.56: Organ Dry ER and Reverberation
The dry signal along with early reflections and reverberation.

Track 24.57: Organ Dry and Reverberation
The dry signal and reverberation together, excluding the early reflections. Notice the distinction between the dry sound and its reverb, whereas the previous track has a much more unified impression and firmer spatial localization.

Plugin: Sonnox Oxford Reverb

Reverberation (late reflections)

The reason that the term *reverberation* is used in this text and not *late reflections* is that the term *reflection* suggests something distinct, like an echo. However, the reflection pattern succeeding the early reflections is so dense that it can be regarded as a bulk of sound. Sometimes, reverberation is referred to as the *reverb tail*. The reverberation consists of reflections bounced from many surfaces many times [path (c) in Figure 24.6]. As sound is absorbed every time it encounters a surface, later reflections are absorbed more as they encounter a growing number of surfaces. This results in reverberation with decaying amplitude. The **level** of the reverberation is an important factor in our perception of depth and will be explained shortly.

Track 24.58: Organ Reverberation Only
The isolated reverberation used in the previous tracks.

Plugin: Sonnox Oxford Reverb

Pure reflections

Early studies on reverbs have shown that, in comparison to the reverberation, the spaced nature of the early reflections gives our brain different psychoacoustic information regarding the space. Many reverb designers have based their designs on this observation with many units having different engines (or algorithms) for each of these alleged components. Later studies demonstrated that such an assumption is incorrect and that it is wrong to separate any reverb into any discrete components. Another argument commonly put forth is that a room is a single reverb generator and any distinction made within a reverb emulator might yield less natural results. Such an approach was adopted by respectable companies such as Lexicon and Quantec and can be witnessed in the form of the superb reverb simulations that both of these companies are known for.

Reverb ratios and depth

Reverbs are the main tool we employ when adding depth to a mix. In order to understand how this can be best employed, it is important to first understand what happens in nature.

The inverse square law defines how sound drops in amplitude in relation to the distance it travels. For example, if the sound 1 m away from a speaker is 60 dBSPL, the sound 5 m away from the speaker will drop by 14 dB (to 46 dBSPL), and the sound 10 m away from a speaker will drop by 20 dB (to 40 dBSPL). It should be clear that the farther away the listener is from the sound source, the lower in level the direct sound will be. But the direct sound is only one instance of the original sound emitted from the source. The reverberation is a highly dense collection of reflections, and although each reflection has dropped in level, summing all of them together results in a relatively loud sound. Also, as sound bounces repeatedly off walls, it is as if the surfaces of all the walls in a room turn into one big sound source emitting reverb; under such conditions, regardless of where the listener is in the room, the reverb will be roughly the same level, as moving away from one wall will bring us closer to another. For simplicity's sake, we can regard a specific room as having reverberation at a constant level of 46 dBSPL.

If a listener in such a room stands 1 m away from the speaker, they will hear the direct sound at 60 dBSPL and the reverberation at 43 dBSPL. The farther away the listener gets from the speaker, the quieter the direct sound becomes, but the reverberation level remains the same. Put another way, the farther away the listener is, the lower the ratio is between the direct sound and the reverberation. At 5 m away from the speaker, the listener will hear the direct sound and the reverberation at equal levels. We call the distance at which such equality happens *critical distance*. Beyond the critical distance, the reverberation will be louder than the direct sound, which should damage intelligibility and clarity. As the direct sound will still be the first instance to arrive at the listener, it will still be used by the brain to decode the distance from the source, but up to a limited distance.
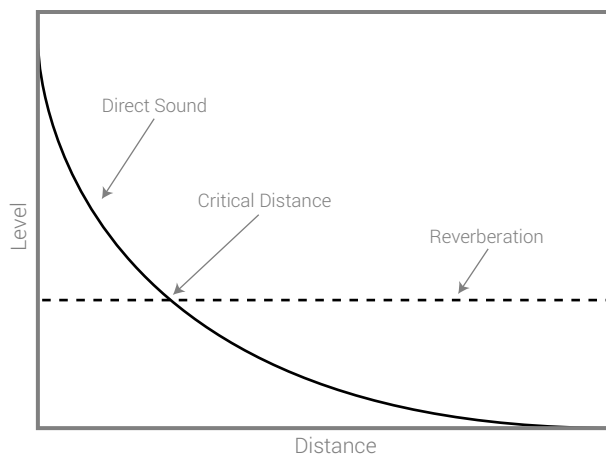
Figure 24.9 Critical distance. The relative levels of the direct sound and the reverberation in relation to the distance between the source and the listener. The distance at which the direct and reflected reverberations are equal in level is known as *critical distance*.

The direct-to-reverberant ratio is commonly used by recording engineers, especially those working with orchestras, to determine the perceived depth of the sound source captured by a stereo microphone pair. For mixing purposes, we decrease the level ratio between the dry signal and the reverb order to place instruments farther away in the mix. Just like in nature, intelligibility and definition can suffer if the reverb is louder than the dry signal. It should be added that the perceived loudness of the reverb is also dependent on its decay time and its density.

Depth positioning during mixdown is commonly achieved by different level ratios between the dry sound and the reverb.

Track 24.59: Depth Demo
This track is similar to Track 6.11, where the lead is foremost, the congas are behind it, and the flute-like synth is farther back. The lead is dry, while the congas and the flute are sent at different levels to the same reverb.

Track 24.60: Depth Demo Lead Down
Compared to the previous track, the lead is lower in level here. This creates a very confusing depth field: On the one hand, the dry lead suggests a frontal positioning, and, on the other, the louder congas, despite their reverb, suggest that they are at the front. When asked which one is foremost, the lead or the congas, some people would find it hard to tell.

Track 24.61: Depth Demo Flute Up
Compared to the previous track, the flute is higher in level. As the reverb send is post-fader, the flute's reverb level rises as well. The random ratio between the dry tracks and their reverb creates a highly distorted depth image.

Plugin: Audio Ease Altiverb
Percussion: Toontrack EZdrummer

Decay time

How long does it take a reverb to disappear? In acoustics, a measurement called RT60 is used, which measures the time it takes reverb in a room to decay by 60 dB. In practical terms, 60 dB is the difference between a very loud sound and one that is barely audible. This measurement is also used in reverb emulators to determine the "length" of the reverb. Scientifically speaking, the decay time should be measured in relation to the direct sound, but some emulators reference it to the level of the first early reflection. In nature, a small, absorbent room will have a decay time of around 200 ms, while a huge arena can have a decay time of approximately 4 seconds.

Decay time hints at the **size** of a room—bigger rooms have longer decay times as the distance between surfaces is bigger and it takes more time for the reflections to diminish.

Decay time also gives us an idea of how reflective the **materials** in the room are—painted tiles reflect more sound energy compared with glass wool.

The decay time in a digital emulator is largely determined by the size of the room. One should expect a longer decay time in a church program than in a small room. Longer decay creates a heavier ambiance, while a shorter decay is associated with tightness. Longer decay also means a louder reverb that will cause more **masking** and possible intelligibility issues. With vocals, there is a chance that the reverb of one word will overlap the next word. This is more profound in busy-fast mixes where there are small time gaps between sounds. If too much reverb is applied in a mix and if decays are too long, we say that the mix is "washed" with reverb.

In many cases, especially when reverbs are used on percussive instruments, the decay time should make **rhythmical sense**. For example, it might make sense to have a snare reverb dying out before the next kick. This will "clear the stage" for the kick and also reduce possible masking. There is very little point delving into time calculations here—the ear is a better musical judge than any calculator, especially when it comes to rhythmical feel. In addition, reverbs may become inaudible long before they truly diminish.

Snare reverbs are commonly automated so they correspond to the mix density; a common trick is to have a shorter snare reverb during the chorus, where the mix calls for more power. As reverbs tend to soften sounds, it might be appropriate to have a longer reverb during the verse. Some engineers have also automated snare reverbs in relation to the note values played—a longer decay for quarter-notes and a shorter decay (or no reverb at all) for sixteenth-notes.

It is worth remembering that a longer pre-delay would normally result in a later reverb decay—lengthening the pre-delay and shortening the decay can result in an overall more present reverb. Level-wise, many find that a short, loud decay is more effective than a long and quiet one.

> *Long decay times can increase masking and damage clarity. When used as a special effect, mainly on percussive instruments, rhythmical judgment must be used.*

Track 24.62: Decay Source
The untreated track used in the following samples.

Track 24.63: Decay Long
A long decay that results in a very present effect that can shift the attention from the beat (this can be desirable sometimes).

Track 24.64: Decay Beat
A shorter decay time that becomes inaudible with the next kick. Note how, in comparison to the previous track, the kick has more presence.

Track 24.65: Decay Lag
Slightly shorter decay than in the previous track; however, this decay makes less sense rhythmically, as the audible gap between the reverb tail and the kick can create the impression that the kick is late and therefore slightly offbeat.

Room size

The room size parameter determines the dimensions of the simulated room, and in most cases it is linked to the decay time and the early reflection pattern. Coarse changes to this parameter distinguish between small rooms such as bathrooms and large spaces such as basketball arenas. Generally speaking, the smaller the room is, the more coloration occurs. Increasing the room size can result in a more vigorous early reflections pattern and longer pre-delay—combined with shorter decay time, the resultant reverb could be more pronounced.

Density

A density parameter on a reverb emulator can exist for the early reflections alone, for the reverberation alone, or as a unified control for both.

The density of the early reflections gives us an idea of the size of the room, where denser reflections suggest a smaller room (as sound quickly reflects and re-reflects from nearby surfaces). The density parameter determines how many discrete echoes constitute the early reflections pattern, and with low values discrete echoes can be clearly discerned. Reducing the early reflections density can minimize the comb filtering caused by phase interaction with the direct sound.

With both early reflections and reverberation densities, higher settings result in a thicker sound that can smooth the sharp transients of percussive instruments. Low density for percussive instruments can cause an unwanted metallic effect similar to flutter echo. But the same low-density settings can retain clarity when applied to sustained sounds such as pads or vocals (which naturally fill the gaps between the sparse reverb echoes). The density setting also relates to the masking effect—a denser reverb will mask more brutally than a sparse one.

Low-density settings can reduce both timbre distortion and masking. High density is usually favorable for percussive sounds.

Diffusion

The term *diffusion* is used to describe the scattering of sound. A properly diffused sound field will benefit from more uniform frequency response and other acoustic qualities that make reverbs more pleasant. Diffusers are commonly used in control and live rooms in order to achieve such behavior. Diffusion is determined by many factors; for instance, some **materials** such as bricks diffuse sound more than other materials such as flat metal. An **irregularly shaped** room also creates a more diffused sound field compared with a simple cubical room. When diffusion occurs, the reflection pattern becomes more complex, both in terms of spacing and levels.

Different manufacturers try to imitate a diffused behavior in different ways. In most cases, the implementation of this parameter is very basic compared with what happens in nature. Many link the diffusion control to density, sometimes in such a way that high diffusion settings result in growing density over time or produce less regular reflection spacing. Density and diffusion are commonly confused as their effect can be identical. The variety of implementations makes it essential to refer to the emulator's manual in order to see exactly what this parameter does. Listening to what it does is even more informative.

Track 24.69: Diffusion Low
With low diffusion, the reverb decay is solid and consistent.

Track 24.70: Diffusion High
High diffusion results in a decay that shows some variations in level and stronger differences between the left and right channels.

Plugin: Cubase RoomWorks
Snare: Toontrack EZdrummer

Frequencies and damping

Frequency treatment can happen at three points along the reverb signal path:

- **Pre-reverb**—where we usually remove unwanted frequencies that can impair the reverb output.
- **Damping**—frequency treatment within the reverb algorithm, which relates to the natural properties of the simulated space.
- **Post-reverb**—where we usually EQ the reverb output in order to fit it into the mix.

Some emulators enable frequency treatment at all three points. Others offer damping only, which can be used as an alternative to post-reverb EQ. If neither pre- nor post-EQ processing are available, a dedicated EQ can always be inserted manually into the signal path—within an audio sequencer, this simply involves inserting an EQ plugin before or after the reverb (on the aux track). In the analog domain, this requires patching an EQ between the aux send output and the reverb. Patching an EQ after the reverb is seldom required as most effect returns include some form of frequency control.

Pre-reverb equalization usually involves pass or shelving filters. Low-frequency content can produce a long, boomy reverb sound that can clutter and muddy the mix. An HPF placed before the reverb can prevent this by filtering low-frequency content—like that of a kick in a drum mix. Some high-frequency content might produce a luminous and unpleasant reverb tail—a pre-reverb shelving EQ can resolve such a problem.

Damping is concerned with the behavior of frequency over time. High frequencies are easily absorbed: it takes 3″ (76 mm) of acoustic foam to eliminate all frequencies above 940 Hz; it takes a bass trap 3′ (1 m) deep to treat a frequency of 94 Hz. High frequencies are also absorbed by the air itself, especially when traveling long distances in large spaces. The natural reverb of an absorbent space has its high frequencies decaying much faster than low frequencies, resulting in less high-frequency content over time.

One could be mistaken for thinking that a pre- or post-shelving EQ could achieve an identical effect, but such a static treatment is found mostly within cheap reverb emulators. A better simulation is achieved by placing filters within the emulator's delay network, which results in the desired attenuation of frequencies over time.

The damping parameter usually represents the ratio between the reverb decay time and the frequency decay time. More specifically, a decay time of 4 seconds and an HF damping ratio of 0.5 means that the high frequencies will decay within 2 seconds. Damping ratio can also be higher than 1, in which case specific frequencies will decay more slowly over time. Further control over damping is sometimes offered in the form of an adjustable crossover frequency. If an emulator has a single control labeled "damping," it is most likely to be an HF damping control (as these are more commonly used than LF damping controls).

Many digital emulators produce brighter reverbs than those produced by real spaces. **Damping HF** can attain more natural results or help in simulating a room with many absorptive materials. But HF damping can be useful in other scenarios—the sibilance of vocals can linger on a reverb tail, making it more obvious and even disturbing (although often mixing engineers intentionally leave a controlled amount of these lingering high frequencies). Noises caused by synthesized sounds that include FM or ring modulation, recordings that capture air movements like those of a trumpet blow, distortion of any kind, harsh cymbals, or even the attack of an acoustic guitar can all linger on a long reverb tail and add unwanted dirt to the mix. In such cases, HF damping can serve as a remedy.

HF damping can result in more natural simulation and also prevent
unwanted noises from lingering on the reverb tail.

LF damping can be applied in order to simulate rooms with materials that absorb more low frequencies than high frequencies, such as wood. Sometimes low-frequency reverberation is required, but only in order to give an initial impact. Employing LF damping in such cases will thin the reverb over time and will prevent continuous muddying of the mix.

Post-reverb equalization helps to tune the reverb into the mix. After all, a reverb is a long sound that occupies frequency space that might interfere with other signals. Just as relevant is a discussion about how different frequencies can modify our reverb perception. High frequencies give the reverb presence and sparkle, which

many choose to retain, especially when reverbs are used for an unnatural effect. Conversely, a more transparent, even hidden reverb can be obtained if its high frequencies are softened. Very often, high frequencies are attenuated in order to create a warm or dark effect—such a reverb can be heard on many mellow guitar solos. Attenuating high frequencies can also result in an apparent increase in distance. Low frequencies make reverbs bigger, more impressive, and warmer. A boost with a low-shelving EQ can accent these effects. The more low frequencies a reverb has, the bigger the space will appear. LF attenuation (or filtering) will thin the reverb and make it less imposing.

> *The high-frequency content of a reverb contributes to its presence, while the low frequencies contribute to its thrill.*

Track 24.71: RevEQ Untreated
The source track for the following samples involves both the dry and wet signals. Note the sibilance lingering on the reverb tail.

Track 24.72: RevEQ Pre-Shelving Down
A shelving EQ at 1.25 kHz with –6 dB of gain is placed before the reverb. This eliminates much of the lingering sibilance, but results in a muddy reverb that lacks spark.

Track 24.73: RevEQ Damping HF Down
The reverb's damping band mid-hi crossover frequency was set to 1.25 kHz, and the damping factor was 0.5. The resultant effect is brighter, and the lingering sibilance has been reduced drastically. Compared to the untreated track, the high frequencies on this track decay much faster.

Track 24.74: RevEQ Post-Shelving Down
The same shelving EQ as in Track 24.72 is placed after the reverb. Despite sounding dark like Track 24.72, it is possible to hear the lingering sibilance being attenuated by the post-reverb filter.

Track 24.75: RevEQ Damping HF Up
With the damping factor set to 1.5, it is easy to hear how high frequencies take much longer to decay.

Track 24.76: RevEQ Damping LF LMF Down
By damping the low and low-mid frequencies by a factor of 0.5, the reverb retains its warmth despite losing some sustained lows.

Plugins: t.c. electronic VSS3, PSP MasterQ

Reverbs and stereo

Mono reverbs

A certain practice in drum recording involves a single room microphone in addition to the stereo overheads. Although the room microphone is positioned farther away from the drum kit, and therefore captures more reverb compared with the overheads, it is

Tools

interesting to hear how the stereo recording from the overheads can sound more reverberant, spacious, and natural.

There are a few aspects that contribute to the limited realism of mono reverbs. One of them is the masking effect, where the direct sound masks the early reflections. An early trick to solve this involved panning the reverb away from the direct sound. But, while this can yield an interesting effect, it sounds unrealistic. The main issue with a mono reverb is that it does not reassemble the directional behavior of a natural reverb, which arrives at our ears from all directions. A reflection pattern that hits the listener from different angles will sound more realistic than one that arrives from a single position in space. Although our pair of monitors can usually only simulate one-sixth of the space around us (and this even happens on one dimension only), a substantial improvement is achieved by using a stereo reverb.

Stereophonic behavior is essential for the reproduction of a lifelike reverb.

Track 24.77: Percussion (Dry)

Track 24.78: Percussion (Stereo Reverb)

Track 24.79: Percussion (Mono Reverb)
The dry source is monophonic. Note how the stereo reverb produces a realistic sense of space while the mono reverb is more reminiscent of a retro effect.

Plugin: Trillium Lane Labs TL Space
Percussion: Toontrack EZdrummer

Having said that, it might be surprising to learn that mono reverbs are not rare in mixing. For one, they are used as a distinctive effect. Then, they are used when some properties of a stereo reverb are unwanted, notably its width. To give an example, say we have two distorted guitars, each panned hard to a different extreme, and we want to add a sense of depth to these guitars or just add a bit of life. Sending both guitars to a stereo reverb will fill the stereo panorama with harmonically rich material that might decrease the definition of other instruments. In addition, the localization of the two guitars could blur. Instead of sending both guitars to a stereo reverb, we can send each to a mono reverb and pan both the dry and wet signals to the same position. The mono reverbs might not produce a natural space simulation, but they would still provide some sense of depth and life without filling the stereo panorama. It is worth remembering that usually using a single channel from a stereo reverb will translate better than panning both channels to the same position; this is due to the phase differences that many reverbs have between the two channels.

Mono reverbs are used either as a special effect or when the width of a stereo reverb is unwanted.

Track 24.80: Guitars (Dry)
The dry guitars panned hard-left and hard-right.

Track 24.81: Guitars (Stereo Reverb)
Each guitar sent to a different extreme of the same stereo reverb.

> **Track 24.82: Guitars (Two Mono Reverbs)**
> Each guitar is sent to a mono reverb, which is panned hard to the respective extreme. It is worth noting that there is a more realistic sense of ambiance with the stereo reverb, although the two mono reverbs still provide a strong spatial impression. Also, compared to the stereo reverb, the two mono reverbs can appear somewhat tidier.
>
> The following tracks are identical to the previous three, but only involve one guitar:
>
> **Track 24.83: Guitar (Dry)**
>
> **Track 24.84: Guitar (Stereo Reverb)**
>
> **Track 24.85: Guitar (Mono Reverb)**
>
> Plugin: Trillium Lane Labs TL Space

Stereo reverbs and true stereo reverbs

The original digital reverb proposed by Schroeder was capable of mono-in/mono-out. A decade later, Michael Gerzon designed a stereo-in/stereo-out reverb. Needless to say, stereo-in reverbs support the common configuration of mono-in/stereo-out, but, as many hardware units will use different algorithms for mono and stereo inputs, an exact selection of input mode needs to be made. Software plugins are self-configuring in that respect.

Not all the reverbs that accept a stereo input process the left and right channels individually. Some reverb emulators sum a stereo signal to mono either at the input or at a specific stage such as after the early reflections engine. If a reverb sums to mono internally, panning a stereo send might not result in the respective position imaging—the generated reverb can remain identical no matter how the feeding signal is panned. An emulator that maintains stereo processing throughout is known as a *true stereo reverb*.

The implementation of a stereo reverb requires different reflection content on the left and right output channels. It is common for the two channels not to be phase-coherent, which is why soloing a reverb will often make a phase meter jump to its negative side. These phase differences result in a spacious reverb that spans beyond the physical speakers and can seem to be coming from around us. The drawback of these phase differences is their flimsy mono-compatibility. Some reverbs give different controls that dictate the overall stereo strength—the more stereophonic a reverb is, the more spacious it will appear.

Panning stereo reverbs

The practice of panning a stereo reverb return to the extremes is not always justified. When panned in such a way, the reflections will occupy the full width of the stereo panorama. The potential masking that a reverb might cause can be reduced if a narrower panning tactic is used. There are many cases where narrowing a stereo reverb is desirable. For example, a strong stereo reverb on a snare can sound foreign to the mix if panned to the extremes. More cohesive results can be achieved if the reverb is panned around the source instrument so it only occupies a slice of the stereo

panorama. Identical panning schemes can be applied to nearly every reverb that is used as an effect, including vocal reverbs.

In less natural mixes, even the ambiance can be narrowed down. This creates a less spacious impression, but a more intense effect. It also keeps the stereo extremes clear for other sounds. As narrowing a reverb makes it more monophonic, it can result in a decrease in perceived depth, which can work well in a powerful mix. However, if the ambiance is narrowed too much, it can give the impression that the instruments are located in a long tunnel.

Although the narrowing of the reverb output is commonly done using pan-pots, these can cause unwanted phase interaction between the left and right channels. It is worth checking whether the emulator has a stereo spread control (sometimes called stereo width). With some emulators, using the stereo spread control instead of panning the reverb return can minimize the phase interaction between the two channels and produce a healthier reverb altogether.

If reverbs are used as a special effect or to fill stereo gaps, they do not have to be panned mirrored around the center. For instance, if a snare is panned to 11:00, it is reasonable to pan the snare reverb between 10:00 and 12:00 (Figure 24.10).

A stereo reverb need not be panned to the stereo extremes or be placed symmetrically around the center.
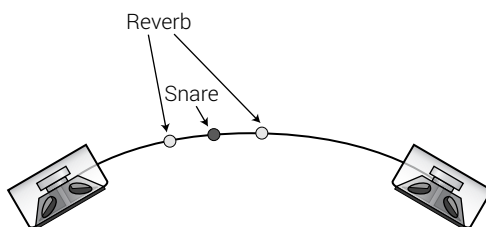
Figure 24.10 Possible snare reverb panning.

Track 24.86: Reverb Panning Full Width
The stereo reverb in this track is panned to the extremes; therefore, it spans the full stereo panorama.

Track 24.87: Reverb Panning 6 Hours
The reverb is panned between 9:00 and 15:00. Notice that there is still a good sense of space while the extremes are clear.

Track 24.88: Reverb Panning 3 Hours
The reverb is panned between 10:30 and 13:30. Its narrow image is reminiscent of the sound at the opening of a tunnel.

Track 24.89: Reverb Panning Mono
Both reverb channels are panned center, resulting in mono reverb. This yields an unnatural impression, and the dry signal masks some of the reverb qualities.

Plugin: t.c. electronic MegaReverb

Track 24.90: Panning SnareVerb Source
The source track used in the following samples.

Track 24.91: Panning SnareVerb Full Width
Although sometimes desirable, such a wide panning of the reverb results in an unnatural spatial image. It feels as if the drums are in one room and the snare is in another.

Track 24.92: Panning SnareVerb Half Width
Narrowing the reverb around the snare makes it sound more like an effect than a spatial reverb.

Track 24.93: Reverb Panning Mono
Panning both extremes of the stereo reverb to the center results in a mono reverb. Even more than in the previous track, the reverb appears as an effect here, but the lost stereo width has altered the overall impression.

Plugin: Universal Audio Plate 140
Drums: Toontrack EZdrummer

Other reverb types

Fixed-source reverb emulators

Most digital reverb emulators can be seen as a virtual room in which artificial reverb is simulated. Despite this, most reverb emulators do not offer the ability to position the sound source (or the listener) in the virtual room. This can be down to, among other issues, processing power limitations. Many of today's emulators are still based on the principles of the original Schroeder design. For simplicity, we shall regard such emulators as having a fixed position for both the sound source and the listener, and we will therefore call them "fixed-source" reverbs.

Figure 24.11 illustrates the reflections of a real room where a violin and a trumpet are placed in front of a stereo microphone pair. The very first reflection to arrive at the left microphone will be that of the trumpet, and it will be immediately followed by the first reflection from the violin.

When we aim to simulate this recording scenario with a fixed-source emulator, we feed it with a stereo mix of both the trumpet and the violin. To reflect their position in Figure 24.11, the trumpet will be panned slightly to the left and the violin will be panned center. Both channels of the stereo mix will contain both the trumpet and the violin, although at different level ratios.

Figure 24.12 illustrates what will happen within the reverb emulator, where the speakers represent the stereo input signal and the microphones represent the stereo output. In this case, the first early reflection will consist of a mix between the trumpet and the violin; in fact, all reflections will be made of such a mix between the two instruments.
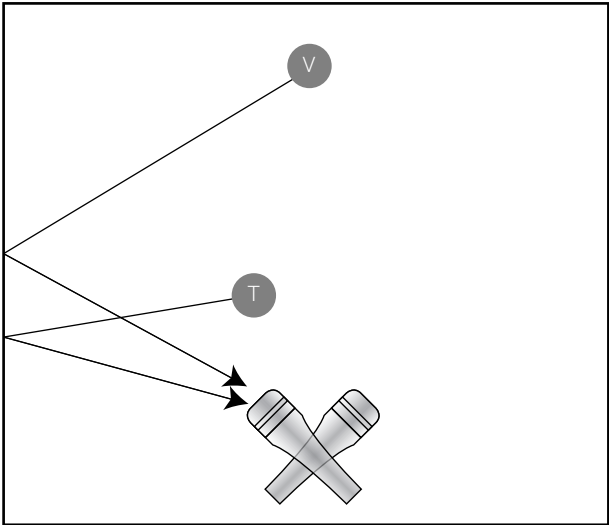
Figure 24.11 The very first early reflections in a real room from a violin (V) and a trumpet (T).
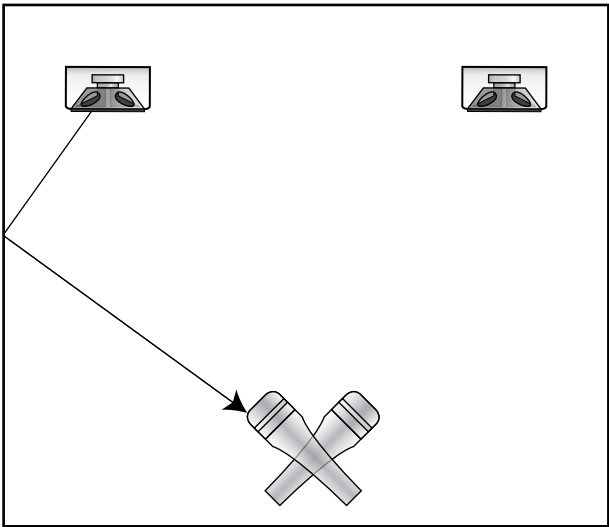
Figure 24.12 Early reflections reproduction within a fixed-source reverb emulator.

In a scenario like the one in Figure 24.12, which reminds us very much of the way sound is recorded in reverb chambers, lies one of the great limitations in feeding any reverb emulator with a composite stereo mix. Reverb emulators are incapable of separating the instruments in a stereo mix, and therefore cannot produce an individual reflection pattern per source. When all reflections originate from such a stereo mix, our perception links the relative position of the different instruments. Accurate localization and depth can be achieved if each of the instruments has a distinct reflection pattern that is not correlated to any other instruments. Even tiny reflection variations between the different instruments can make a big difference.

One way to overcome this, usually only possible in the digital domain, is to send each instrument to a single emulator. Each of the emulators has to be loaded with an identical program, with slight variations in parameters between the emulators. The small variations are likely to create a distinct reflection pattern for each instrument, but unlikely to cause ambiance collision. Nevertheless, the fixed position of the source in the virtual room can still limit the extent to which realistic localization can be achieved.

Movable source reverb emulators

Movable source reverbs (sometimes just *source reverbs*) give us control over the position of the sound source in the virtual room. Control over stereo sources is achieved by front–back positioning of either the imaginary speakers or listener; a mono signal can also be moved left and right. The Sonnox *Oxford Reverb* (Figure 24.13) is such a plugin, enabling source positioning via its position control.

Positioning of multiple instruments in one virtual space requires sending each of the instruments to a different emulator with an identical program loaded. The only parameter that will change between each emulator is the position of the source (or the listener). Such a setup will bring about an individual early reflection pattern for each instrument and therefore much more faithful localization and more defined depth.
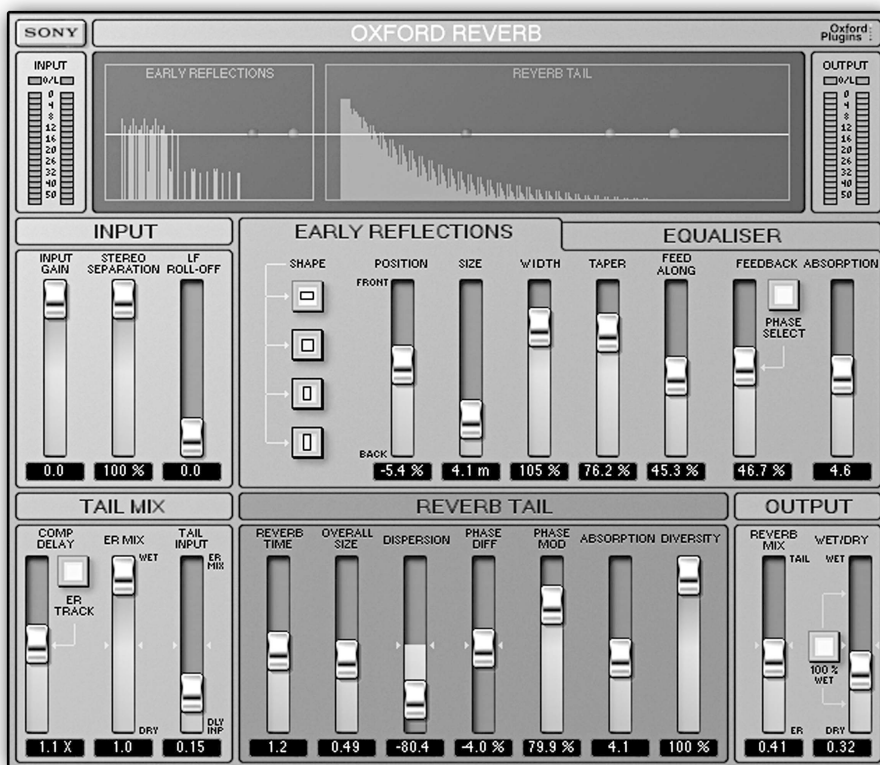
Figure 24.13 The Sonnox *Oxford Reverb*. The position of the source within the space algorithm is adjustable via the position control.

While some digital reverb emulators enable flexible positioning, some convolution reverbs simply involve different impulse responses taken at different positions within the room. The distance is often mentioned in the preset name (for example, 1 m away, 3 m away, and so forth). Some convolution reverbs even provide a graphic display of the room and show the different positions at which the impulse response was recorded.

Gated and nonlinear reverbs

The reflections in a real room can build up and decay in amplitude very differently from those produced by a reverb emulator. While most emulators' loudest reflection is the first one, the reflection pattern in some spaces can build up for a while before reaching maximum amplitude—an equivalent to the attack stage on a dynamic envelope. In addition, the reflections in a real room fluctuate in level, unlike the *linearly* decreasing reflection levels in some reverb emulators. The shape of the reverb envelope is important for natural simulation, and some emulators provide related controls. Mixing engineers can employ various tools to make the decay of reverbs far less natural. Despite potentially damaging its natural impression, such a nonlinear reverb can have more impact.

The decay of a reverb is what we most commonly shape. The simplest way of shaping it is by inserting a compressor after the reverb. Often the compressor is set to moderate threshold, moderate ratio, and fast attack. For truly wild decay shapes, the release can be set to minimum so some pumping takes place; for more subtle results, the release time can be lengthened.

The most famous nonlinear reverb is the *gated reverb*. Rumor has it that it was discovered by accident during a recording session of Peter Gabriel's third album. It is mostly recognized for its use in Phil Collins' "In the Air Tonight." In its most simple form, gated reverb is achieved by inserting a gate after the reverb and setting the gate threshold so it opens with the initial reverb burst. Usually a fast attack and release are used, along with hold settings that make rhythmical sense. This configuration does not alter the shape of the decay curve until its later stages—during the gate hold period, the reverb decay remains as before (Figure 24.14b). One rationale for gating a reverb in such a way has to do with the fact that with sustained dry signals, the reverb is likely to
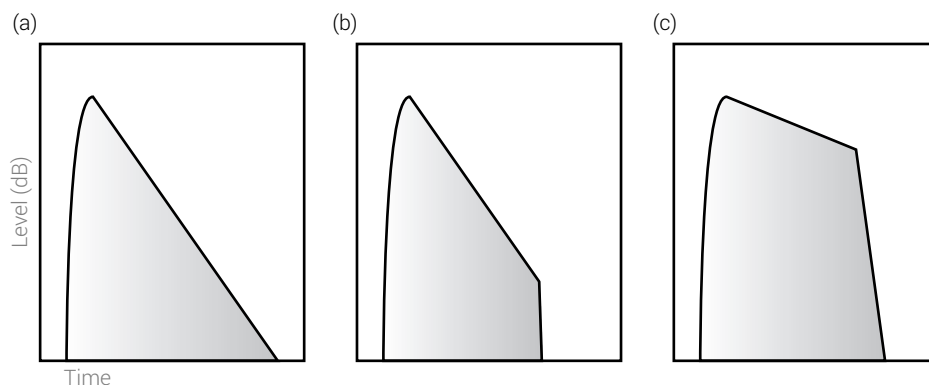
Figure 24.14 Reverb decay. (a) A linear reverb decay. (b) A simple gated reverb decay. (c) A decay that first falls slowly and then dives quickly.
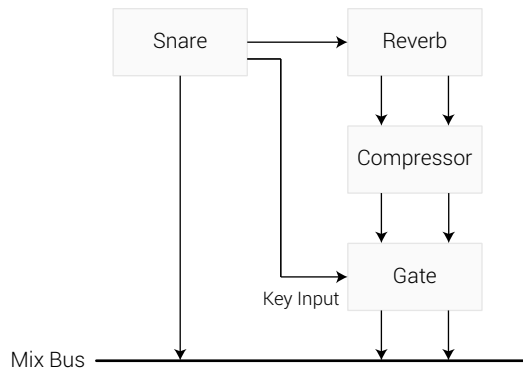
Figure 24.15 Gated snare setup. A schematic signal flow for an "improved" gated reverb.

be masked by the direct sound as soon as it drops by 15 dB. In a busy mix, the reverb tail can be masked by many other sounds—a gated reverb can sharpen the sonic image by cutting these inaudible parts of the reverb.

If instead of decaying linearly the decay shape is altered so it sustains at maximum level (or falls slowly) and then decays abruptly, the reverb becomes more prominent. To achieve this, a compressor is inserted after the reverb. The compressor is set to flatten the decay so it hovers longer at higher amplitudes. But this will result in a reverb that might not drop below the gate's threshold. Therefore, the original sound is sent to the gate's key input. Figure 24.15 demonstrates such a setup.

There are endless numbers of productions in which gated reverb has been applied on snares and toms, but it can be applied to many other instruments, even those of a non-percussive nature. Shaping the reverb's decay envelope can give the reverb more punch or more defined rhythmical characteristics (which must be observed while tweaking the gate). Gated reverbs are so popular that many digital reverb emulators include such a program. Still, better control is usually achieved by setting up the independent building blocks.

While gated reverbs are often used on drums, nonlinear reverbs in general can give the reverb more entity and a stronger rhythmic form.

Track 24.94: SnareTom Source
The source drum sequence before the addition of reverb.

Track 24.95: SnareTom Reverb Only
Both the snare and the toms are sent to the same reverb, which on this track is neither compressed nor gated.

Track 24.96: SnareTom Compressed Reverb
The reverb in this track has been compressed but not gated.

Track 24.97: SnareTom Compressed Gated Reverb
This is essentially the common gated reverb effect. The reverb is both compressed and gated.

Track 24.98: SnareTom Uncompressed Gated Reverb
Removing the compressor preceding the gate results in the effect heard on this track.

Plugin s Sonnox Oxford Reverb, Digidesign DigiRack Expander/Gate Dyn 3
Drums: Toontrack EZdrummer

Reverse reverbs and preverb

A reversed reverb does not really reverse any audio, but it does reverse the amplitude envelope of the reverb, so it starts quietly, grows gradually, and then cuts abruptly (Figure 24.16a). This effect gives a backward impression that can be used as a special effect.
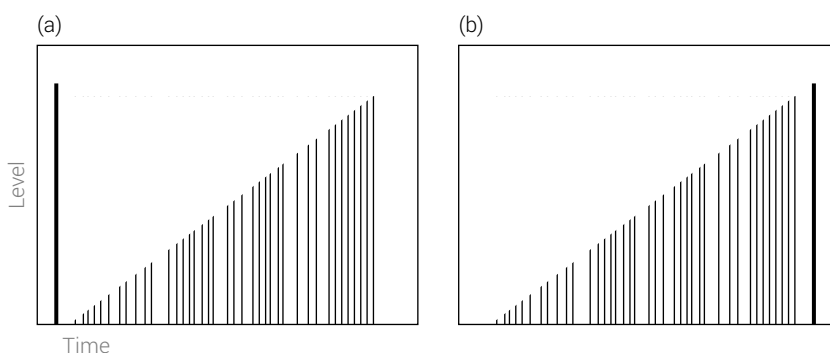
Figure 24.16 (a) A reversed reverb. (b) Preverb. The bold line denotes the direct sound.

A more common practice in mixing is called preverb. This is done by recording the reverb and then playing it backward so that it rises before the actual sound (Figure 24.16b). In order for this to work, the reverb has to be fed with a backward version of the original material. With tapes, this is easily done by reversing the reels, recording the reverb on to an available track and then reversing the reels again. In the digital domain, the original material has to be reversed, the reverb has to be bounced, and then both the original sound and the reverb have to be reversed again.

Track 24.99: Reversed Reverb Mono
Both the dry voice and the reverb are panned hard-center in this track.

Track 24.100: Reversed Reverb Stereo
The voice is panned hard-left and the reverb to hard-right. This makes the reversed development of the reverb easier to discern.

Plugin: Logic EnVerb

As with the previous two tracks, the following two demonstrate preverb.

Track 24.101: Preverb Mono Track 24.102: Preverb Stereo

> **Track 24.103: Preverb Before Reversing**
> This is the voice and the reverb, before reversal.
>
> Plugin: Universal Audio RealVerb Pro

Reverbs in practice

Reverb quantity and allocation

A common mixing discussion relates to the amount of reverbs one should use in a mix. To start the debate, one should imagine oneself in a standard professional studio back in 1971. Such a studio had mostly analog equipment, maybe up to three high-range dedicated reverb units, perhaps along with a couple of less expensive units and a few multi-effect units that had reverb programs as well. It would be fair to say that for most mixes in such a setup, the high-range units were allocated to ambiance and vocals, while other reverbs were allocated based on their quality and the importance of the instruments sent to them. Any mix that involved more than eight reverbs would be considered reverb-happy, with many mixes completed with no more than five reverbs and some with as few as two. Previously, there was little choice other than to feed the same reverb unit with more than one track using an aux send. While there are mixes from that era that involve more than a dozen reverb emulators, it is worth remembering that many beloved mixes from the past were mixed with just a few reverb units. This practice still remains.

The inherent problem with hardware units is that they cannot be multiplied with a click of a mouse like plugins can. Even in 2001, many mixing engineers had to limit the number of reverb plugins due to processing power limitations, and reverbs had to be connected in parallel and in many cases were fed with a mix of many tracks using an aux send.

It is only in recent years that computers have become powerful enough to accommodate even 100 reverb plugins in a single mix. A different reverb plugin can be used as an insert on each individual track and one can also feed a single vocal track into five reverb plugins connected in series. Nowadays, for many of those who software-mix, having eight reverbs in a mix does not seem lavish at all.

If mixing engineers back in 1971 had had an array of 20 reverb units, they might have used them all. But two issues come to mind. The first is the fact that it is easier to grow a tree than a forest. Put another way, a mixing engineer might be tempted to use many reverb plugins while compromising on quality and giving less attention to each reverb. Selecting and configuring a reverb can be a truly time-consuming business, especially for the less experienced engineers who are not yet familiar with their tools. Second, the use of too many reverbs, especially if a different program is loaded onto each, can cause what is known as *ambiance collision*—a blend of too many spaces, which does not make any spatial sense. A certain objective when using reverbs is the creation of a firm, convincing ambiance even when an imaginary ambiance is involved. Although many reverbs can work together, a certain risk exists if insufficient attention is given to

how they interact. Generally speaking, when using more than one ambiance reverb, softening the highs of the short reverbs can help conceal them and reduce collision.

While there is almost total agreement that expensive emulators are superior in producing a natural reverb, it is sometimes the cheapest and most basic emulators that can produce a distinctive effect. One advantage that comes with software plugins is that there is an ever-growing selection—some of today's plugins offer unique controls and sound that might not be found in expensive emulators.

Track 24.104: MultiAmb Dry
The dry signal used in the following samples.

Track 24.105: MultiAmb One Reverb
The woodblock, tambourine, and congas are all sent to the same reverb emulator.

Track 24.106: MultiAmb Three Reverbs
The three instruments are each sent to a different reverb emulator loaded with a different preset. We can still localize the instruments in the depth field, but the simulated space is less coherent than in the previous track.

Plugin: Universal Audio DreamVerb
Percussion: Toontrack EZdrummer

Aux send or insert?

Convention has it that reverbs are fed via auxiliary sends. The most obvious advantage in using a send is that we can send more than one signal to the reverb emulator, and therefore save processing power or share high-quality emulators. Connecting reverbs as a send on an analog desk requires just a few more easy steps compared with an insert connection. However, in the digital domain, the same task usually requires a little extra work; for example, the creation of auxiliary tracks and bus routing. If many reverbs are used, but each is fed from one single track, the multitude of auxiliary tracks can clutter the mixer, and many buses are employed purely for routing purposes (rather than summing). Therefore, some argue that in specific situations, it is sensible to use an insert rather than a send, for instance when a reverb is added to a kick for two bars only, or when the reverb only adds a touch of early reflections that are meant to enrich a mono signal. One thing to remember is that when connecting as an insert, the reverb should output both the dry and wet signals; on both analog and DAW setups, changing the amount of reverb requires more effort than altering the send level, like when reverbs are connected as sends.

Another important point worth considering is that both in the analog and digital signal flows, the inserts are pre-fader. Thus, connecting reverbs as an insert forces a pre-fader feed that leaves all dry/wet adjustment to the reverb emulator. Accessing the emulator every time we want to change the dry/wet balance is somewhat less convenient compared with accessing the send controls.

In all but very specific situations, reverbs are added using sends.

Pre- or post-fader?

As just mentioned, reverbs connected as an insert are pre-fader (with rare exceptions). When connected as an aux send, we have a choice between either a pre- or post-fader send. A post-fader send is also post-cut. Similarly, a pre-fader send is also pre-cut.

The basic principle behind using post-fader send is that the level of the dry signal and the reverb level are linked—when the dry signal level increases, the reverb level increases as well. In practice, sometimes a fader adjustment might not lead to the desired reverb level change, and some additional adjustment to the reverb level will be needed. A post-fader send is necessary if any level or mute automation should take place—when we fade out or mute the dry signal, we want the reverb to fade out or mute as well.

Another point relates to the common use of destructive solos during mixdown. Say we have vocal and guitar both sent to an ambiance reverb. The reverb is solo safe to prevent it from being cut when other instruments are soloed. If the guitar send is pre-fader and the vocal is soloed, the guitar reverb will still be heard along with the vocal and its reverb. This can be highly annoying. Post-fader sends ensure that such situations do not occur.

> *Post-fader reverb send is most common in mixing, and is requisite if any level or mute automation is to take place on the source signal.*

In some situations, we need to bounce the reverb separately from its dry signal. In many software applications, this can be achieved by soloing the reverb auxiliary track, but, as software solos are often destructive, this will cut the dry signal track as well. In order to still have the dry signal feeding the auxiliary bus, and consequently the reverb unit, a pre-fader send can be used. Another situation where a pre-fader send can be beneficial is when the original track is mixed at a very low level, and therefore the post-fader send will not provide enough level to produce an audible reverb.

Pre-compressor reverb send

Some interesting results can be achieved if the dry signal is compressed but the signal sent to the reverb emulator is taken before the compression takes place. In a performance with vibrant dynamics, increased expression is conveyed by louder phrases. Such a performance might be a problem in a mix and require compression. But, if sent to a reverb unit before being compressed, the dynamics of the performance will trigger more reverberation during the louder moments. This "manual automation" effect is most suitable on a performance with quick dynamic changes, such as some vocal performances.

While on an analog desk a copy of the uncompressed signal can always be taken from the multitrack return or a pre-processing insertion point, software users might scratch their heads wondering how this can be done—a software send (the reverb) is always post-insert (the compressor). Figure 24.17 illustrates how this can be done in Logic—the dry signal is sent to a bus, on which the reverb is inserted. Note that the vocal track output is not set to the main mix bus, but to another bus on which the compressor is inserted.

Figure 24.17 A pre-compressor reverb send in Logic.

Track 24.107: Post-Comp Rev Send
The uncompressed vocal used in the following samples.

Track 24.108: Post-Comp Rev Send
The reverb send in this track succeeds the vocal compressor.

Track 24.109: Pre-Comp Rev Send
In this track, first the uncompressed vocal is sent to the reverb emulator, only then being compressed. Notice how the reverb gets louder in relation to the level of the uncompressed voice.

How loud? How much?

Back in the 1980s, when the use of digital reverb units became widespread, much excitement surrounded this new tool and reverbs manifested in many mixes. As time went by, mixing engineers started to make more relaxed use of reverbs, and the general tendency today is that reverbs serve the mix rather than being a sound in their own right. In some of today's mixes, a trained ear will be able to identify many more reverbs than initially seem present—reverbs nowadays tend to be more transparent and subtler.

This might go against the intuition of many novice mixing engineers, who, having introduced it, wish to hear the reverb in the mix. But the beauty of a reverb can persist even if it is not very obvious. One way to check the effectiveness of reverbs, even if hidden, is by muting them and seeing how they contribute to the mix. Yet, in specific situations, having a bold, well-defined reverb is our aim; solo saxophone in a mellow production could be one of these situations.

Tuning ambiance

The standard way to create ambiance involves sending various instruments to one or more ambiance reverbs. However, the resultant ambiance can suffer from various issues due to the frequency content of the various instruments sent to the reverb. Playing such a crucial part, the ambiance in a mix can be enhanced using a specific advanced technique. It involves sending a modified version of the instruments sent to the mix bus to the ambiance reverbs. On an audio sequencer, this can be achieved by duplicating a specific track and disabling its standard output routing to the main mix. Yet, this track is sent to the ambiance reverb instead of the original. We can then process (mainly equalize) the ghost track in different ways, so as to tune it to the ambiance created by the reverb it is sent to.

Delay and reverb coupling

A very common mixing technique involves a blend between a delay and a reverb. Most often, this technique is used on vocals, but it can be applied to other instruments just as well. Blending a delay with a reverb is known to result in a more impressive effect than using one or the other. One way to look at it is as if the delay enhances the reverb. Multitap delays with two active taps (each panned to a different side of the panorama) are highly suitable for this task. Most often, the delay and the reverb are connected in parallel (i.e., neither feeds the other). On an audio sequencer, this can be achieved by sending a track to a bus, then feeding the bus into two different aux tracks—on one a reverb is loaded, on the other a delay. It also pays to try to connect a delay before a reverb—such an arrangement can also produce an appealing effect.

Track 24.110: Vocal Reverb No Delay
This track involves a reverb, but no delay.

Track 24.111: Vocal Reverb With Delay
The delay in this track is added in parallel to the reverb. Notice how, compared to the previous track, the effect here sounds slightly livelier and richer.

Plugins: Audio Ease Altiverb, PSP 82

Reverb coupling

Sometimes we combine two reverbs in order to achieve a specific task. However, this has to be done with caution or some problems might arise. *Coupled reverb* is the term used to describe a "reverb of a reverb." In real life, such a phenomenon can occur when a reverb from one space travels to an adjacent space that adds its own reverb. Coupled reverbs are mainly a concern for natural reverb simulations—they place a space within a space, they increase the reverb time, and they introduce reverbs that can have broken decay slopes. In mixing, coupled reverbs can be introduced with the addition of a reverb to a recording that already contains a reverb or when connecting two reverb emulators in series.

Recordings that include a reverb, unless initially intended, can sometimes create great difficulties for the mixing engineer, and in extreme situations can be unusable. Such recordings are sometimes done by inexperienced recording engineers who may choose to record a guitar amp or an electronic keyboard with their internal (sometimes monophonic) cheap reverb effect. Sometimes a recorded reverb is simply the result of a very reverberant recording environment. One way or another, such reverbs are then imprinted into the recording, often sound cheap, and leave the mixing engineer between a rock and a hard place. It is worth remembering that reverbs often become louder after being compressed. Compressing a vocal recording that contains reverb can result in solid vocal levels but fluctuating reverb. While it is only possible to achieve a completely reverb-free recording in an anechoic chamber, recording engineers can achieve a direct-to-reverb ratio that will present no limitations during mixdown.

> *With the exception of a recording that was intended to capture ambiance, a reverb on a recording is like oil on water—they are likely to not mix well.*

Nonetheless, a combination of two reverb emulators during mixdown can be useful. A valid practice is to send the same signal to two reverb emulators that are connected in parallel. This creates a composed reverb that can have a richer texture. Another possible experiment that can yield an interesting effect is using two mono reverbs and panning each to an opposite extreme.

Yet another trick relates to the way emulators work—many of them have separate engines for the early and late reflections. Some emulators feed the early reflections into the late reflection engine, while others have complete separation between the two. Often we can choose the balance between the early and late reflections or simply switch off one engine or the other. This yields some interesting possibilities where we can combine two emulators, each producing either the early or later reflections. If we have two emulators and one is better at producing early reflections while the other is better at producing reverberation, we can mute the weak engine in each emulator to achieve a reverb that has the best of both worlds. This can work exceptionally well when a convolution reverb is the generator of the early reflections. Convolution reverbs tend to produce a highly believable early reflection pattern, but might be very CPU-consuming if the long reverberation is produced as well. By assigning the convolution reverb with the task of producing the early reflections only, and assigning the less critical task of late reflections to an algorithmic reverb, we can save CPU power, yet end up with more convincing results. The emulators can be connected either in parallel or in series, depending on whether we want to feed the early reflections into the late ones (Figure 24.18).
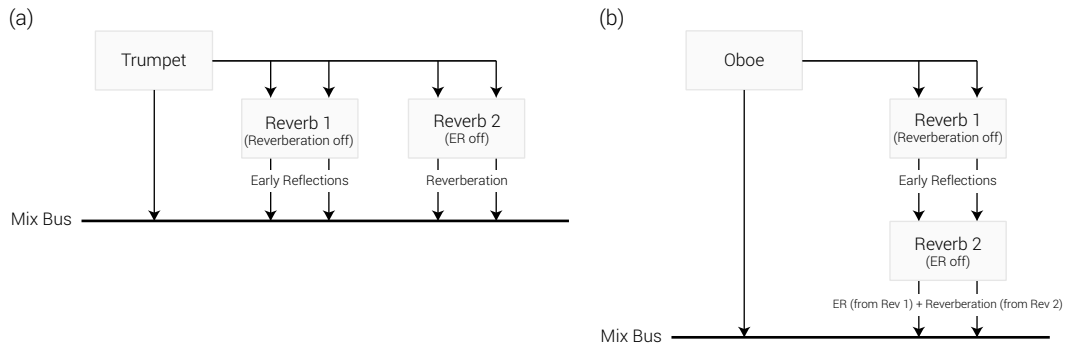
Figure 24.18 Reverb coupling. (a) Parallel reverb setup where each emulator produces a different reverb component. (b) Serial reverb setup where the early reflection from the first emulator feeds the second emulator, which only generates reverberation.

25 Distortion

Background

Mixing is a dynamic art. Mixing trends change through the years, and throughout past decades the extent to which and the way we use different tools has changed as well. There is little argument that distortion is the tool that has gained the most popularity in recent years. Distortion, in that specific mixing context, is not exactly the screaming pedal we use with guitars, but the more subtle harmonic distortion, in its various forms.

There are two main reasons why distortion has become so widely used. First, pop mixes nowadays tend to be more aggressive. This applies to both the analog and digital domains, where distortion is added to various degrees depending on the production. Strictly speaking, the use of distortion is not limited to rock or metal productions—even genres such as dance, chart-pop, and trip-hop might involve a certain degree of distortion. Second, distortion is used to compensate for what can be described as the "boringly clean" digital sound—the inherent distortion of analog equipment and tape media is not an asset of digital audio. Engineers use distortion to add some edge to digital audio by introducing some degree of "appealing error."

Distortion basics

Harmonic distortion and inter-modulation

In the recognition of sounds, our ears use both the fundamental and its harmonics. Harmonics, as already explained, are integer multiples of the fundamental, and thus the two are closely related. By emphasizing harmonics, we increase the definition of sounds. Lower harmonics are louder than higher harmonics. For example, the second harmonic is louder than the 20th harmonic. Due to their relatively loud level, the few low-order harmonics have some extra importance. The second harmonic is an octave above the fundamental, and the fourth harmonic is two octaves above the fundamental. Being octave-related to the fundamental, both harmonics are considered musical—emphasizing them tends to give appealing results. The third and fifth harmonics are also important, but neither is octave-related to the fundamental—both tend to produce a slightly more colorful sound.

A linear device is one that has perfect proportionality between input and output. The transfer curve of such a device is one straight line. An example of a nonlinear device is a compressor with any ratio other than 1:1. When a signal passes through a nonlinear system, different kinds of distortion are produced. The less linear the system, the more profound the distortion. One type of distortion is *harmonic distortion*—essentially, added harmonics. Analog components are incapable of being perfectly linear. A specification called *total harmonic distortion* (THD) measures the harmonic distortion content produced by an analog device under standard test conditions. There are different flavors to analog distortion. The ratio between the low-order harmonics produced by a tube is different from that produced by a transistor. This is a major contributor to the different sounds that a tube and solid-state equipment produce. Although technically speaking the lower the distortion the better, harmonic distortion is an intimate part of the analog sound in general and the characteristics of analog gear in particular. Digital systems are capable of being perfectly linear and thus might not produce harmonic distortion. Although it is technically superior, many find the digital sound lifeless and pale.

Another type of distortion is *inter-modulation*. Like total harmonic distortion, it can be measured and the specification given is simply called inter-modulation (IMD). Just like harmonic distortion, inter-modulation distortion involves additional frequencies, but unlike harmonic distortion these are not necessarily harmonically related to the sound. Therefore, inter-modulation is often considered harsh and it is mostly unwanted. Yet, it is an inseparable part of any nonlinear system.

The problem with distortion in the digital domain

Often distortion produced within an analog system is more musical than that produced within a digital system. The reason for this originates from the fact that analog systems do not have a frequency limit like digital systems do. The highest frequency a digital system can accommodate is called the *Nyquist frequency*, and it is always half the sample rate. The problem with digital systems is that any content produced within the system that exceeds the Nyquist frequency mirrors around it. For example, say we have a system running at a sample rate of 44,100 Hz. The highest frequency such a system can accommodate is 22,050 Hz (the Nyquist frequency). Then, say an 8-kHz sine wave is distorted. The resultant harmonic distortion would include a second harmonic at 16 kHz and a fourth harmonic at 32 kHz—both have an octave relationship with the fundamental. However, since the 32-kHz harmonic is higher than the Nyquist frequency, it will mirror around it, resulting in an alias frequency at 12.1 kHz. This frequency is not harmonically related to the 8-kHz fundamental and will produce a harsh sound. On the same basis, any distortion content that exceeds the Nyquist frequency creates aliasing frequencies below it. On an analog system, no such thing happens—any distortion content above 20 kHz is simply unheard.

We can minimize the harsh aliasing phenomenon of digital processing by using higher sample rates. If, for example, the sample rate of the system above was 88.2 kHz, the Nyquist frequency would be 44.1 kHz. This means that the 32-kHz harmonic would not alias—it simply would not be heard. Even content that does alias might do so into the inaudible range between 20 and 44.1 kHz. With 88.2 kHz, only content above 68.2 kHz would alias back into the audible frequency range. The problem with using higher sample rates is that more samples have to be processed. Compared to a sample rate of

44.1 kHz, an 88.2-kHz audio will require twice the samples processed, which essentially halves the processing capabilities of the CPU (or any other digital processor, for that matter). In practice, the actual processing overhead in using higher sample rates can be more than initially expected—some plugi-ns will need to perform substantially more calculations in order to operate at higher sampling rates.

High-quality plugin developers take this aliasing problem into account in their designs by implementing internal upsampling and then downsampling. Essentially, even if the project's sample rate is 44.1 kHz, the plugin upsamples the audio to, say, 176.4 kHz (4×). This results in a new Nyquist frequency of 88.2 kHz, and only content above 150.4 kHz will alias back into the audible frequency range. The plugin then performs its processing, which might produce content above the original Nyquist frequency—22.05 kHz. Then any content above 20 kHz is removed using an anti-aliasing filter (essentially a high-quality LPF) and the audio is downsampled back to the original project's sample rate. This way, any content that would otherwise alias back into the audible range is simply removed.

Track 25.1: Hats Upsampling
These highly distorted hi-hats were produced while the plugin ×16 upsampling option was enabled.

Track 25.2: Hats No Upsampling
The added frequencies in this track are the result of aliasing frequencies. This track was produced while the upsampling option was disabled.

Plugin: bram@SmartElectronix.com Cyanide 2

It should be mentioned that the aliasing distortion produced within a digital system is often inaudible. The problem is not as bad as one might think. Some manufacturers do not perform upsampling for this reason. Still, accumulating distortion (a distortion of distortion) could become audible.

Parallel distortion

Just like the parallel compression technique, where a compressed version is layered underneath the original, signals are sometimes distorted and then layered below the original. This gives us added control over the amount of distortion being added. Consequently, this lets us drive the distortion harder to produce a stronger effect, but then layer it underneath at lower levels so it is not too obvious or crude.

The following tracks demonstrate parallel distortion on vocal and drums:

Track 25.3: Vocal Source
The unprocessed source track.

> **Track 25.4: Vocal Distortion Layer**
> The vocal in the previous track distorted. In the following tracks, the unprocessed track and this distorted layer are mixed, with the distorted layer a few dB below the unprocessed track.
>
> **Track 25.5: Vocal Distortion –18 dB**
>
> **Track 25.6: Vocal Distortion –12 dB**
>
> **Track 25.7: Vocal Distortion –6 dB**
>
> **Track 25.8: Vocal Distortion 0 dB**
> In the following tracks, the same principle as before is demonstrated on drums:
>
> **Track 25.9: Drums Source**
>
> **Track 25.10: Drums Distortion Layer**
>
> **Track 25.11: Drums Distortion –26 dB**
>
> **Track 25.12: Drums Distortion –18 dB**
>
> Plugin: SansAmp PSA-1
> Drums: Toontrack EZdrummer

Ways to generate distortion

Gain controls

As already described in Chapter 10, a boost on the gain control, affecting the signal early in the signal path, can overload successive components. In the analog domain, we usually speak first of *overloading*, which mostly adds subtle distortion, and only then *saturation*, which can produce drastic distortion. The nature of solid-state components is that they provide increasing distortion with gain. Tube equipment tends to produce a very appealing distortion up to a specific point, where the system seems to "break" and produce very undesired clicks. Tapes can be overloaded using the same principle by boosting the input signal. Bipolar transistors, FET, tube, and tape distortion all have different qualities as they all produce different harmonic content. Also, the more each system overloads, the more compression occurs, which makes the overall effect even more appealing.

Digital clipping

Of all the types of distortion, digital clipping is probably the least appealing. The harsh limiting of audio exceeding the 0-dB threshold on a digital system can produce an extremely unpleasant sound, especially if the added content aliases. There are also different types of digital clipping, where differences are dependent on what happens to signals that exceed the system limit. On most floating-point systems, these signals are trimmed to the highest possible value (simply hard-limiting). However, different integer notations can produce different results, where exceeding signals can alias around the
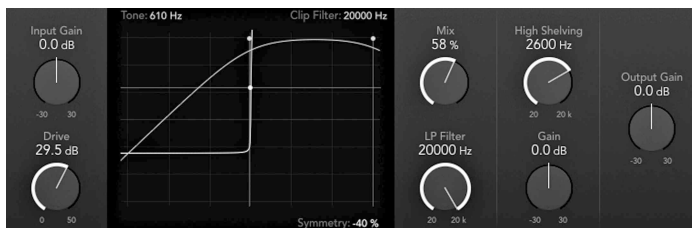
Figure 25.1 The Logic *Clip Distortion* plugin.

highest sample value; they can be displaced to the bottom of the value range or displaced across the zero line.

The following tracks demonstrate clip distortion. The dB in the track names denote the approximate number of dB by which the signal exceeded the 0-dB limit of the system.

Track 25.13: Vocal 6-dB Clip Distortion

Track 25.14: Vocal 18-dB Clip Distortion

Track 25.15: Drums 6-dB Clip Distortion

Track 25.16: Drums 18-dB Clip Distortion

The above tracks are produced by overloading the mix bus in Pro Tools LE. (After bouncing, the 18-dB versions are attenuated as a measure of ear safety.)

The following tracks are produced using Logic Clip Distortion. The drive is set to 30 dB, and the symmetry to 0 or 60 percent (versions 1 and 2, respectively):

Track 25.17: Vocal Clip Distortion 1

Track 25.18: Vocal Clip Distortion 2

Track 25.19: Drums Clip Distortion 1

Track 25.20: Drums Clip Distortion 2

Plugin: Logic Clip Distortion
Drums: Toontrack EZdrummer

Generally speaking, clipping tends to produce the strongest type of distortion. Some plugins try to imitate the clipping sound of analog devices by introducing additional processing along with the basic digital clipping. The clip distortion in Figure 25.1 is one of them. It is worth noting the "Mix" control, which lets us blend the undistorted and distorted sounds.

Short attack and release on dynamic range processors

In Chapter 17, which covered compressors, we saw that a short attack and release can act within the cycle of low frequencies, thus altering the shape of the waveform and producing distortion (refer to Figure 17.29 for a demonstration of this). This type of

distortion is not reserved solely for compressors and can be produced by short time constants on any dynamic range processor. We can use this type of distortion to add some warmth and definition to instruments with low-frequency content, notably basses.

Wave shapers

A wave shaper is a pure implementation of a transfer curve. Essentially, the signal passes through a transfer characteristics function over which we have some control. This is somewhat similar to a dynamic range processor with no time constants and an unrestricted transfer curve. Wave shapers are not very common and might be a sub-facility within a different effect. However, they can produce extremely interesting results that can be used in subtle amounts for gentle enhancements or in a drastic form for creative effect.
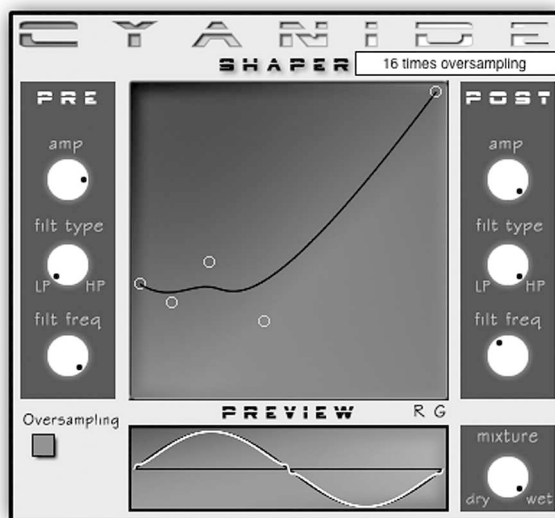
Figure 25.2 The Smartelectronix *Cyanide 2* plugin (donationware).

A few tracks demonstrating the impressive distortion capabilities of wave shapers:

Track 25.21: Wave Shaper 1 (Drums)

Track 25.22: Wave Shaper 1 (Vocal)

Track 25.23: Wave Shaper 2 (Drums)

Track 25.24: Wave Shaper 2 (Vocal)

Track 25.25: Wave Shaper 3 (Drums)

Track 25.26: Wave Shaper 3 (Vocal)

Track 25.27: Wave Shaper 4 (Drums)

Track 25.28: Wave Shaper 4 (Vocal)

Track 25.29: Wave Shaper 5 (Drums)

Track 25.30: Wave Shaper 5 (Vocal)

Plugin: bram@SmartElectronix.com Cyanide 2
Drums: Toontrack EZdrummer

Bit reduction

Early digital samplers that emerged around the early 1980s had low specs compared with those used today—in addition to low sample rates, they were designed around 8-bit (and later 12-bit) samples. Even now, the 8-bit sound is much sought after in genres such as hip-hop, where many find the lo-fi sound of drum samples appealing.

When we discussed software mixers in Chapter 11, we discussed the importance of dither, which rectifies the distortion caused by bit reduction. Unsurprisingly, like many other areas in mixing where we are after the "technically wrong" and the less precise sound, bit reduction can also have an appealing effect. Essentially, a bit reduction process simply reduces the bit depth of digital audio. Although within the audio sequencer the audio is still represented by 32-bit float numbers, sample values are quantized to the steps of the target bit depth. For example, reduction to 1 bit would produce 32-bit float values of 1.0 and –1.0 only (in practice, however, most processors also output 0.0 as a possible sample value). This process produces quantization distortion, where the lower the target bit depth is, the more the distortion. Clearly, we wish to keep this distortion, so no dither is applied. Just like with drum samples, bit reduction can be used to give a lo-fi sense to various instruments. It can also be used as a creative effect.
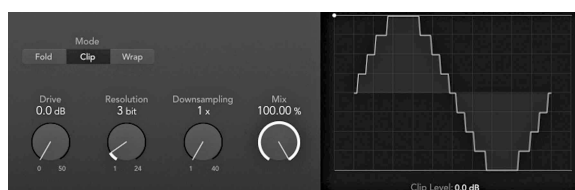
Figure 25.3 The Logic *BitCrusher*. This plugin combines bit reduction, three-mode clipping distortion, and digital downsampling. In this screenshot, only the reduction to 3 bits is utilized. The eight quantization steps that would affect a sine waveform are visible on the display.

Track 25.31: Moroder Source
The source track, before applying bit reduction.

Track 25.32: Moroder 5 Bits
All instruments but the kick are reduced to 5 bits in this track. While the lows of the bassline become wimpy, the effect on both the hats and the snare is usable.

Plugin: Logic BitCrusher

Amp simulators

A famous mixing practice in the analog domain is called *re-amping*—feeding a guitar recording during mixdown back into a guitar amp so a different sound texture and distortion can be created. Mostly, the signal sent to the amps during mixdown is the direct recording (before any pedals or amp). Many plugins nowadays are designed to imitate the sound of classic guitar cabinets, and usually these plugins also include other guitar processors such as tremolos, compressors, gates, echo engines, and guitar reverbs. When there is a direct recording of a guitar, we can easily choose the final guitar sound during mixdown (together with the cabinet recording, perhaps replacing it).

Even before the digital age, re-amping was not limited to guitars. Sometimes engineers sent vocals, drums, or other tracks into a guitar amplifier (or guitar pedals), using the amp as a distortion or effects unit. Using the same principle, an amp simulator plugin can be used to distort any type of material. Like any other type of distortion, subtle settings can be useful for gentle enhancements and more trashy settings for more drastic results.

Figure 25.4 The Universal Audio *Nigel*. The center blocks on this plugin are the amp simulator where different amps and cabinets can be chosen. The side blocks provide additional effects such as tremolo and delay.

Track 25.33: dGtr SM57
A microphone recording of a distorted guitar.

Track 25.34: dGtr Direct
This is the direct recording (before amp or pedals) of the same recording in the previous track.

Each set in the following samples contains three tracks. The first is the direct recording from the previous track processed with an amp simulator. The second and third are vocal and drums being processed with the same preset.

Track 25.35: Amp Simulator 1 (Guitar)

Track 25.36: Amp Simulator 1 (Vocal)

Track 25.37: Amp Simulator 1 (Drums)

Track 25.38: Amp Simulator 2 (Guitar)

Track 25.39: Amp Simulator 2 (Vocal)

Track 25.40: Amp Simulator 2 (Drums)

Track 25.41: Amp Simulator 3 (Guitar)

Track 25.42: Amp Simulator 3 (Vocal)

Track 25.43: Amp Simulator 3 (Drums)

Track 25.44: Amp Simulator 4 (Guitar)

Track 25.45: Amp Simulator 4 (Vocal)

Track 25.46: Amp Simulator 4 (Drums)

Track 25.47: Amp Simulator 5 (Guitar)

Track 25.48: Amp Simulator 5 (Vocal)

Track 25.49: Amp Simulator 5 (Drums)

Plugin: McDSP Chrome Amp
Drums: Toontrack EZdrummer

Dedicated units or plugins

As distortion becomes more popular, more dedicated units and even more plugins are appearing in the market. The *Culture Vulture* by Thermionic Culture is one respected hardware unit. The Soundtoys *Decapitator*, shown in Figure 25.5, is an analog saturation modeler plugin that can produce a variety of distortion options.

Figure 25.5 Soundtoys *Decapitator* plugin.

26 Drum triggering

A short story. A band goes into a studio to record a song. The producer and the engineer spend some time miking up the drum kit. They try different microphones in different positions and mic each drum with two microphones. They experiment with different snares of different sizes and woods, go through the process of tuning the toms as they wish, build the famous kick tunnel, and dampen all the drums. After perhaps 7 hours of preparation, they cut the first drum take. Two weeks later, the track reaches the mixing engineer, who, in their own studio, at night, replaces all the drum recordings with samples—perhaps 20 minutes of work. None of the original drum tracks are left in the mix. The mixing engineer also filters some ambiance from the overheads and reconstructs it from the drum samples. Three days later, the clients come around to hear the mix. The drums sound amazing. The band is happy, the producer is happy, the record company is happy, and 4 months later the listeners are happy—everyone is happy. Many tracks have similar stories; and to some extent, one of them is Nirvana's "Smells Like Teen Spirit."

The practice of using samples to add to or replace real drums started somewhere around the early 1980s. It was used in pop and soon became an extremely common practice in the more energetic metal genres (where it is still prevalent). Throughout the 1990s, engineers started using it more and more in virtually any genre that involved recorded drums. This technique has always had its critics, who dismissed it as being disloyal to the true essence of performance and recording. Nonetheless, drum triggering is an *extremely* popular practice nowadays, and many recorded productions involve some degree of drum replacement.

One approach to generating the samples is programming a drum sequence that matches the recorded performance, and then using a sampler to trigger the drums. Before audio sequencers existed, this meant syncing a tape machine to a hardware sequencer. Unless the recorded drums were tight to a metronome, this involved (and still does) the generation of tempo maps. Altogether, this method can be rather tedious. Today, the most common approach to drum triggering involves feeding a drum track, say the snare, to a machine or a plugin that triggers the samples in real time. The samples can be those of a drum module, taken from a sample library, or just clean individual hits that were recorded during the session. The samples either **replace** or are **layered over** the original recording. Although most commonly used with drums, we sometimes use samples to reinforce cymbals.

There are various benefits of drum replacement. First, we have total freedom over the sample we use, so if we do not like the **sound** of, say, the toms, we replace them with tom

samples we do like. When layered, samples might be used to add missing timbre components. Second, the resultant samples track is **spill-free**, which makes gating, compression, and the addition of reverb both easier and more consistent. Third, most often as part of the process of drum triggering, we have ultimate control over the level of the hits, which can make the **performance** more consistent (and therefore possibly more powerful).

Drum replacement is ideal for budget drum recordings carried out in problematic rooms. The good news is that you can get away with horrendous close-mic recordings. The bad news is that you still need good overheads recording, which, out of all drum recordings, is the hardest to achieve in problematic rooms—the small dimensions of the room can cause comb filtering, profound room modes can alter the captured frequency content, and, above all, the reverb in such rooms is often not the greatest. Still, drum replacement is a blessing in situations where the drum recordings are flawed. Similar to guitar reamping, it allows us to select the sound of our drums during mixdown.

The drums in the following demonstration are from the "Hero" production.

Track 26.1: Kick Original
A blend between the recorded kick in and out microphones.

Track 26.2: Kick Trigger
The original audio kick track is converted to MIDI using Logic's "Audio to Score" function, and then Toontrack EZdrummer is used to trigger this kick sample. Note the consistent velocity, unlike in the previous track.

Track 26.3: Snare Original
A blend between the recorded snare top and bottom microphones, and a snare reverb.

Track 26.4: Snare Trigger
The snare trigger is a product of Trillium Lane Labs' Drum Rehab. This trigger is not intended to replace the original snare, just to blend in with it.

Track 26.5: Snare Original and Trigger
A blend between the original snare and its trigger. Compare this track to Track 26.3.

Track 26.6: Drums Without Triggers
The recorded drum kit without triggers.

Track 26.7: Drums With Triggers
In this track, the original kick is replaced with the kick trigger, while the snare trigger is only layered with the original snare.

Plugin: Trillium Lane Labs Drum Rehab
Kick: Toontrack EZdrummer

Methods of drum triggering

Delay units with sample and hold

One early method of triggering drum samples was using a delay unit that offered a sample and hold facility. A single drum hit was sampled and then could be triggered by feeding the

unit with the performance itself. The main disadvantages of this method were that it required sampling a drum hit first (rather than choosing between different prerecorded samples) and that the triggered sample was always the same, including its level.

Drum module triggers

Another early method, still used today, involves the drum modules used with electronic drum kits. These modules have trigger inputs that are fed from the drum pads via a standard audio connection. Instead of feeding these inputs with the pads, we can feed a line-level signal from the multitrack or perhaps a gated version from the desk. Drum modules let us choose different samples via their interface. Most modules are velocity-sensitive, meaning that the triggered sample level corresponds to the level of the input hit. Some modules also provide dynamic control, where we can set the ratio between input and output velocities. For example, we might configure the unit, so wild variations in the level of input hits will only result in small (or no) level variations of the output samples.

Footswitch on a synchronizer

Some synchronizers have an input footswitch. The synchronizer can generate a specific MIDI note every time a signal exists at the footswitch input. By routing a drum track into the footswitch input, the synchronizer generates a MIDI note with every hit. We can connect the MIDI output of the synchronizer to a drum module or a sampler.

Drum replacement plugins

Early drum replacement plugins, such as Digidesign's *Sound Replacer*, worked offline. Recent products such as *TL Drum Rehab* (Figure 26.1) and *Drumagog* perform drum replacement in real time. Essentially, these plugins are loaded as an insert on a drum track, various detection parameters are configured, and we get to choose from various drum samples on our hard drive. We can also set the mix between the original track and the triggered samples. Most replacement plugins let us set the ratio between input and output velocities. Also, triggering might be based on velocity layers, where

Figure 26.1 The Trillium Lane Labs *Drum Rehab*.

different input velocities trigger different samples, each of the same drum played at different velocities. This gives more realistic results.

Audio to MIDI

Some applications let us convert audio into MIDI data. In Logic, the single audio track is first analyzed offline, and, after configuring various settings a new MIDI region, is created with notes representing each drum hit (Figure 26.2). The advantage of audio-to-MIDI conversion is that we can then more easily alter the various hits, whether in level or in timing. Digital Performer offers a very similar conversion that takes place in real time via a plugin called *Trigger* (Figure 26.3).
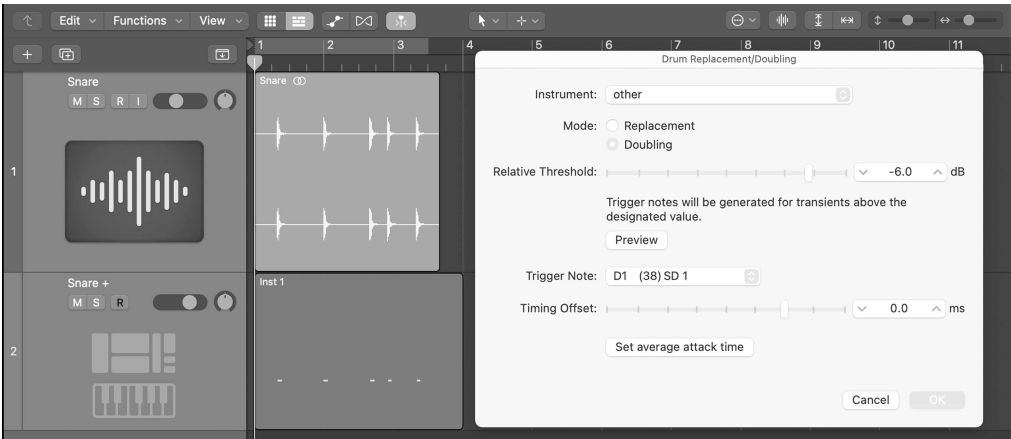
Figure 26.2 Logic's "Audio to Score" conversion. The top track shows the original kick track, the floating window shows the "Audio to Score" dialog with its various configuration options, and the bottom track shows the resultant MIDI track after all notes have been transformed to Cl.

Figure 26.3 The MOTU *Trigger* plugin.

After the audio is converted into MIDI, the track can feed either a sampler or a software instrument. Various software instruments are designed specifically for drums and can be extremely handy for the purpose of drum triggering. One of them is the Toontrack *EZdrummer* shown in Figure 26.4.

Figure 26.4 The Toontrack *EZdrummer*. This multilayer drum sampler has various drum sets; within each, different drums can be selected. This type of product is very useful during drum triggering.

Manually pasting samples onto a track

Although not exactly triggering, on an audio sequencer we can copy a specific drum sample, and then go from hit to hit and paste the sample to replace the original drum hits. This is achievable with toms or crashes, where there are not many hits to replace. However, it can be very tedious with snares and kicks, of which there might be more than 100 hits. Locating the beginning of each hit can also be a difficult task. Features such as Pro Tools' tab-to-transient and beat-detective can be of real help in these situations as they enable you to precisely locate each hit. Then keyboard macros like those offered by applications such as Quick Keys can automate the process so that hits can be replaced in a matter of seconds.

27 Other tools

To cover each and every processor and effect available for mixing engineers would be exhausting. Combined with the huge number of plugins that emerge each year, a whole new book could be written. This chapter covers a few mixing tools that may come in handy and are therefore worth knowing about. Since these are less common than the tools covered so far, the information is kept short.

MS

Background

Mono vinyls utilized lateral (horizontal) groove excursions to encode amplitude variations. When stereo records were introduced in 1958, there needed to be a way to encode both the left and right channels. One proposal was to encode the left channel as lateral excursions and the right channel as vertical ones. The problem was that, had a stereo record been played on a mono player, only the left channel would have been heard (due to the infancy of pan-pots back then, this could have meant losing the drums, vocals, or bass). As part of his stereo invention, Alan Blumlein taught us that any left and right information can be converted into middle and side, and vice versa. Based on this, EMI cut the first stereo record encoding the middle as lateral groove excursions and the side as vertical. This way, even if a stereo record was played on a mono player, the player would only decode the lateral motion, which represented the mid (the mono sum of left and right). A stereo player would decode both lateral and vertical motions (mid and side) and would convert these into left and right. This mid-side system is used in vinyl pressing to this day. Stereo transmissions of FM radio and television are also encoded in MS.

MS (or *mid-side* or *mono/stereo*) is the general name given to stereo information existing in the mid-side form. It is a well-known stereo recording technique that has already been discussed in Chapter 14 (see also Figure 14.15). Many mixing engineers come across MS at least once—the famous *Fairchild 670* could work in either left/right mode or mid/side mode (termed Lat/Vert on the *670* for its relation to records). Some plugins provide MS operation modes in addition to the standard left/right.

Figure 27.1 The Universal Audio *Fairchild* plugin. This plugin looks and works very similarly to the original unit (the price of which can easily exceed the $25,000 mark due to its vintage status). Like the original unit, the plugin can work in either left/right or mid/side modes (determined by the control between the two channels). The mid/side mode is labeled Lat/Vert (lateral/vertical) as per the stereo record cutting method.

MS and LR conversion

The conversion from LR to MS is straightforward: mid is half the sum of left and right, while side is half the difference between left and right. When such a conversion takes place, the mid is sent to the left channel and the side to the right. The conversion from MS to LR is even simpler: left is the sum of mid and side, while right is the difference between them. Expressed in equations:

- $M = (L + R) / 2$
- $S = (L - R) / 2$
- $L = M + S$
- $R = M - S$

We can achieve such conversion with any desk or audio sequencer, but doing so is time-consuming and convoluted. Fortunately, we have plugins to do the conversion for us. Looking at the equations above, the conversion from LR to MS is very similar to that from MS to LR, only that during the LR to MS conversion we have to attenuate by 6 dB (halving). Indeed, the same plugin can perform both conversions, and we might need to manually reduce the MS output by 6 dB. In addition, we can use a plugin that works on only one of two channels. This way, we could first convert from LR to MS, process either the M or S channels, and then convert back to LR. Figure 27.2 demonstrates this.
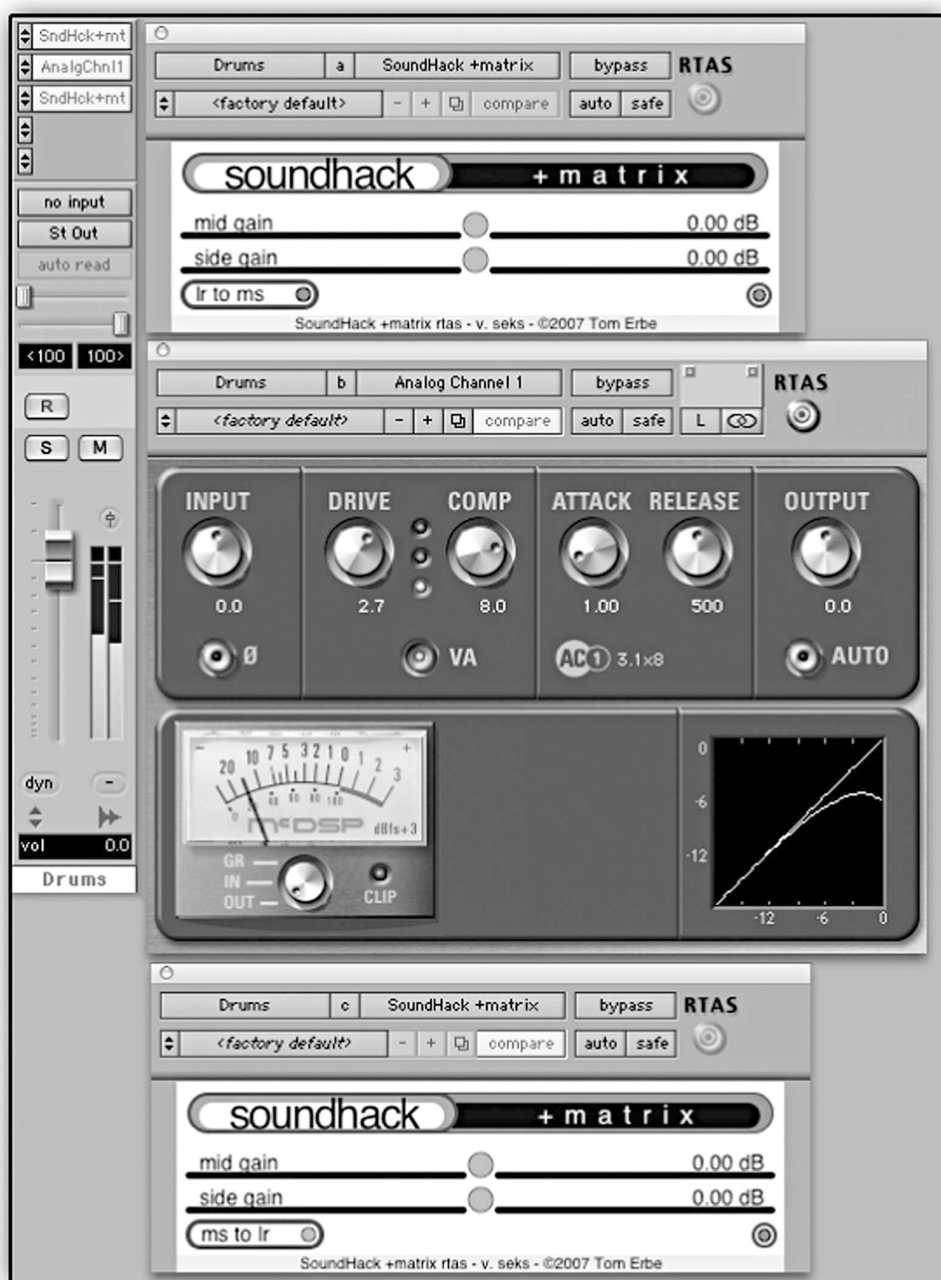
Figure 27.2 MS processing. The +*Matrix* plugin by SoundHack (freeware) is loaded on the first insert slot and converts LR to MS. Pro Tools provides a multimono mode, where a plugin can process the left and right channels individually. The McDSP *Analog Channel 1* is loaded in such a mode on the second insert slot and only affects the left channel, which represents M. Then, on the third insert slot, another +*Matrix* plugin is loaded, this time converting the MS back to LR.

Applications of MS

There are two principal ways in which MS can be used in mixing. We are discussing stereo tracks here, and it can be easier to understand the whole concept if we imagine overheads panned hard to the extremes. The first use of MS involves changing the ratio between M and S. There is little point attenuating S, since this produces similar results to narrowing the stereo image using the pan-pots. But something very interesting happens when we attenuate M. Basically, this results in a wider, more spacious stereo image, which can be useful sometimes, for example, if we want to widen a stereo overheads recording that presents a narrow image. What effectively happens as we attenuate M is this: instruments panned dead center drop in level, and instruments panned slightly off-center start shifting toward the extremes. Again, this can be practical if the snare on the overheads is only slightly off-center and we want to shift it farther sideways to clear more space for other instruments. But a really interesting thing happens with instruments panned to the extremes. As we attenuate M, such instruments start appearing on the opposite extreme out of phase. Essentially, this creates a similar effect to the out-of-speakers trick. But when discussing the out-of-speakers trick, we said that it works best with low frequencies; high frequencies simply tend to appear around the center. So, as M is attenuated, low-frequency instruments panned to the extremes tend to appear as if coming out of speakers, while instruments with high-frequency content tend to shift inward toward the center. An absorbing effect indeed.

The second beneficial way of using MS involves situations where we want to process the mid or side individually. For example, say we want to attenuate the kick on the overhead since it does not blend well with the close-mic. We can use the close-mic to duck the overheads, but this would mean that any hi-hats, toms, or crashes would be ducked as well. Since on overheads the kick is often found around the center, whereas the hi-hat, toms, and crashes are panned off-center, it makes sense to only duck the overheads in the center (mid) and not both the left and right channels. Another example would be a stereo choir recording fighting with the lead vocal. We can achieve better separation by rolling off around, say, 3 kHz on the mids of the choir. Other, more creative, applications might involve an instrument ducking the S of its own ping-pong delay; this way, while the instrument is playing, its echoes would appear monophonic, and when the instrument is not playing, the slow release of the ducker would slowly make the echoes open into stereo. During mastering, de-essing is sometimes done on the M channel only, as vocals are normally panned dead center; this results in the less noticeable ducking of the center only, rather then both the left and right channels. Many more practical and creative examples exist, but the principle should be clear by now.

Track 27.1: Drums Original
These drums are intentionally programmed to have a narrow stereo width. The kick is panned hard-center. The snare is panned slightly to the left, the hi-hats slightly farther to the left from the snare.

Track 27.2: Drums MS Processed
The drums from Track 27.1 are converted to MS, the M channel is attenuated by 2 dB, and the S channel boosted by 4 dB, and then they are converted to LR format again. The result

is drums with a wider stereo image and both the snare and the hi-hats evidently shifted farther to the left.

Track 27.3: Loop Original LR
The original loop, with the kick panned hard-right, the snare hard-center, and the hi-hats hard-right.

Track 27.4: Loop MS
The loop after MS encoding. The snare, which was previously panned hard-center, now only appears on the M channel (left).

Track 27.5: Loop M Only
Only the M channel of the encoded loop.

Track 27.6: Loop S Only
Only the S channel of the encoded loop.

Track 27.7: Attenuating M
In this track, the M channel changes in level prior to decoding back to LR. This track is subdivided into four-bar chunks. Bars 1–4: the original loop. Bars 5–8: the M channel is gradually attenuated to –∞. Bars 9–12: no M, only S (which is essentially the same S on both left and right channels with one channel out of phase). Bars 13–16: the M is gradually boosted back to 0 dB. Bars 17–20: the original loop again.

Plugin: SoundHack +Matrix
Drums: Toontrack EZdrummer

Pitch shifters and harmonizers

A pitch shifter alters the pitch of the input signal without altering its duration. The way it works is very interesting and provides the basis for pitch correction and granular synthesis. If we zoom into the recording of a human voice, we identify repeating patterns that only vary over time. For example, when one sings a long "ahhhh," one produces many repeating patterns within it. In practice, nearly everything we sing is essentially made up of small repeating patterns. The sound of most instruments is produced in exactly the same manner. A pitch shifter identifies these patterns, then shrinks or stretches them into the same time span. For example, shrinking a pattern to half of its length then fitting two shrunken halves into the same space as the original pattern would raise the pitch by an octave. Doing so pattern by pattern would raise the overall input signal by an octave. To pitch something down, half of the original pattern is placed into the space of the original one. The actual process is more complex, but at its core is this rudiment. Also, percussive instruments do not present clear patterns, so pitch shifting them involves slightly different algorithms.

Pitch shifters provide transposition values in the form of semitones, and often cents. Some provide wet/dry control. They can be used to improve double-tracking or ADT. But their real power is in their ability to artificially produce harmonics. Perhaps most commonly, pitch shifters are beneficial with bass guitars. By layering a bass guitar with an artificially produced second harmonic (an octave above), we can increase the guitar

definition. We can also add the third (an octave and perfect fifth or 19 semitones) and fourth (two octaves above) harmonics and alter the level between them to shape the color of the guitar. We can also blend the bass guitar with a version of itself an octave down. This enhances the low-energy power of the guitar. Indeed, this process is used in many bass enhancers.

Harmonizers differ from pitch shifters in that they produce more than one transposed version of the input signal and are configured around musical intervals and chords. As one would expect, harmonizers can enrich a vocal performance, but can also be used with guitars and keyboards.

Track 27.8: Vocal Delay Only
This track involves a blend of the dry voice and an 85-ms delayed version.

Track 27.9: Vocal Delay with Pitch Shift
The delayed version in this track is pitch shifted by 30 cents. This track sounds more like double-tracking, whereas the previous track sounds more like a plain delay.

Plugin: Logic Pitch Shifter 2

Exciters and enhancers

Exciters and *enhancers* are two terms that can be used interchangeably. Both exciters and enhancers have the role of making instruments, mixes, and masters sound better. Better, for that matter, is often described as a combination of cleaner, brighter, closer, wider, livelier, and more defined. There are two questions to ask: How do they do it? and Why *shouldn't* we always use them?

The first aural exciter was produced by Aphex in 1975. Legend has it that the whole concept was discovered by accident, when one channel of a stereo amplifier faultily produced distortion. When added to the properly working channel, this improved the sound. Today, many companies manufacture exciters, and we also have a multitude of dedicated plugins or sometimes an extra facility within other types of plugins. Although the exact way in which each exciter works is a guarded secret, the core principles of aural enhancement are widely known:

• **Addition of harmonics**—by adding a controlled amount of harmonics, sounds tend to sound brighter and clearer. This is somewhat similar to distorting or saturating signals, only that within an exciter it is done in a more controlled fashion.
• **Level enhancements**—an enhancer might employ a process that increases the perceived loudness of signals. Essentially, part of the device involves some degree of loudness maximization.
• **Dynamic equalization**—by dynamically equalizing the input signal, it can appear more solid and less prone to frequency/level imbalance caused by varying notes.
• **Stereo enhancement**—by introducing subtle variations between the left and right channels, instruments might appear to widen, have more life, and become clearer. We have seen this previously when discussing stereo equalization.

- **Phase enhancement**—frequency-dependent phase shifts are said to occur acoustically and are a proven part of signal processing and reproduction. Enhancers employ a process that delays different frequency ranges by different times, and by doing so rectifies phase misalignments. This can increase the definition and clarity of the signal in question.

Figure 27.3 The Noveltech *Character* plugin. The enhancement process involves just three controls: character, target, and one of three modes.

Another look at this list reveals that all, apart from the last, of the techniques have already been described in this book. Exciters and enhancers often involve very few controls. For example, the more basic enhancers only have an amount pot. While this means they can be employed very quickly, this can also be limiting to some extent—we might be able to achieve similar enhancements using standard tools such as equalizers and distortion, with which we have further control over the final result.

In analog studios, there would normally have been no more than one or two enhancers. Mixing engineers had to choose which instrument (if any) called for enhancement the most. In the plugin age, we can enhance every single instrument. So why wouldn't we just enhance all of them? The answer goes back to the basic mixing question: What is wrong with it? If nothing is wrong with an instrument, why should we enhance it? Even if something is wrong, enhancers might not be the solution—perhaps compression is the answer.

If we enhance each instrument in the mix, they will each sound better but not with respect to one another. From a mix perspective, they might all sound equally exciting. Only enhancing one or two instruments creates some contrast that can be beneficial to the mix. In addition, enhancing one instrument might reduce the clarity of another instrument. This is partly due to the fact that enhancers add frequency content in the form of harmonics that can mask other instruments.

Finally, there is another risk in using enhancers, called *frying* or *overcooking*. We find the immediate effect of exciters and enhancers appealing, and we can easily be tempted to drive an enhancer harder than necessary. This can result in sound that is

fatiguing to the ear, and sometimes it is only later that we discover how overuse of exciters makes the mix sound brittle and undefined.

Without a doubt, exciters and enhancers can produce beautiful and magical results, but the points above must be considered when deciding when they should be used and to what extent. As with many other mixing tools, using them sparingly can often yield the most effective results.

Track 27.10: Bass Source
The source bass track used in the following sample, before enhancement.

Track 27.11: Bass Enhanced
Apart from being louder, this track is also slightly brighter and overall presents more vitality.

Plugin: The "Enhance" facility of the Sonnox Oxford Limiter.

Track 27.12: Drums Source
The source drums to be enhanced in the following samples.

Track 27.13: Drum Enhancement 1
Compared to the source track, there is added warmth and low power in this track.

Track 27.14: Drum Enhancement 2
Compared to the source track, this track is brighter and sounds more alive.

Track 27.15: Drum Enhancement 3
This is frying—this track presents distorted frequency content favoring the highs. Listening to this track too loud or for a long time can easily cause ear fatigue.

Plugin: Noveltech Character
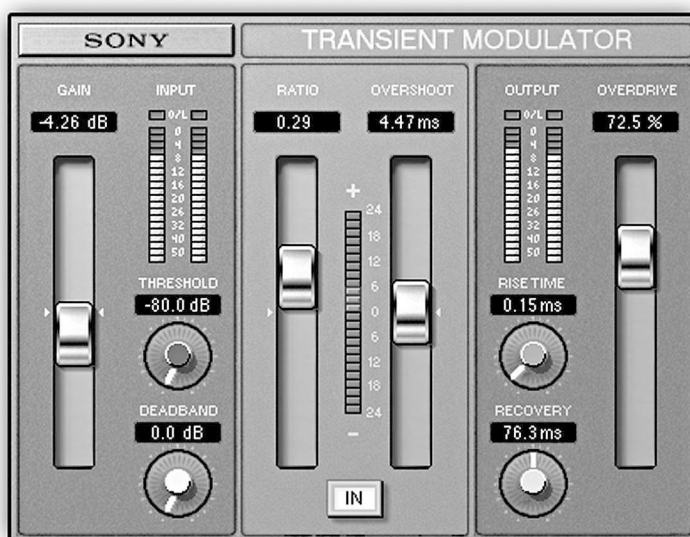
Figure 27.4 The Sonnox *Oxford Transient Modulator* plugin.

Transient designers

Transient designers, or *transient enhancers*, are designed specifically to accent or contain transients. Essentially, most of them are a hybrid of upward expanders (for accenting) and downward compressors (for containing). However, most transient designers examine differences between peak and RMS variation of the input signal in order to determine when transients really happen. Transients rise quickly enough to yield fair differences between the peak and RMS readings; a signal that rises slowly would not. It is for this very mechanism that transient designers have an advantage over upward expanders or compressors—they tend to handle transients more uniformly, regardless of the level at which these happen. Put another way, a quiet snare hit that does not overshoot the threshold would not trigger compression; a transient designer can be set so both quiet and loud hits would be treated.

Like enhancers, transient designers are something of an automatic tool that can easily add punch to percussive instruments. They can also add life and accent or revive the dynamics of various instruments. As with compressors, we can use them to reduce ambiance and decay. Altogether, they tend to produce a very appealing effect, especially if used with more conservative settings.

Track 27.16: Compressed Bass
This is the source track for the following sample. The bass in this track is compressed so as to suppress its attack.

Track 27.17: Bass Transient Revival
After applying a transient modulator, the bass gained some extra attack.

Track 27.18: Drum Compressed
These are the source drums used in the following sample.

Track 27.19: Drum Enhancement 3
After applying a transient modulator, the snare, hats, and the kick have more attack. Note that the ambiance has been reduced and the drum image moved forwards.

Plugin: Sonnox Oxford Transient Modulator
Drums: Toontrack EZdrummer

28 Automation

It is as simple as this—commercial mixes are full of automation. This process, which is performed naturally by professional mixing engineers, is somewhat over-looked by the novice. Even before the invention of multitrack recorders, engineers used to ride levels of different microphones during recording. With the introduction of multitrack recorders, this practice didn't stop, but some of it was postponed to the later mixing stages. Before automation computers were integrated into con-soles, the engineer, the assistant, the producer, and even the band members (as many hands as were needed) used to gather around the console to **perform** automation passes that were printed straight onto the final two-track. Each person knew exactly what needed to be moved, how, and when. For the more complex mixes, an "action score" was written. It was, practically, a performance, and the console was the instrument. If someone botched, the whole performance had to start all over again. Automated consoles were introduced around the late 1970s. Even today, many of them can only write the automation of channel faders, mutes, and solos, and studio engineers and their companions continue to perform live automation on analog desks.

Nowadays, audio sequencers let us write the automation for virtually every control in the mix and, moreover, provide graphical editors in which automation can be corrected or even drawn from scratch. There is no need for 12 hands since we can automate different controls during different passes. Never in the history of mixing was writing and editing automation as easy as it is today with audio sequencers. It is surprising that many DAW users fail to comprehend the benefits of this powerful facility and the potent effect automation can have on the mix.

We say that each song is a story. In the case of classical music, progressive rock, and many jazz pieces, they might be epics. In modern pop productions, the story is often squeezed into around three minutes, and changes happen very quickly. There is a lot happening in each story—the music develops, the arrangement changes, different sections should have different impacts, and the importance of different instruments varies. It could be unfair to have so much action streaming through a static mix. It is our responsibility as mixing engineers to accommodate the dynamic movement of the music and the structural elements of the production. Then, automation can always be used to create some interest or add some extra movement. The options are endless,

493

and virtually any process or effect can be automated. It is possible to regard this late phase in the mix as playtime, where creativity and experiment might replace practical needs. Finally, level automation, mostly on vocals, is sometimes done before a compressor as part of manual balancing.

Any list of possible automation examples would be partial—the options are truly endless. It should be stressed that the most common mix automation (and often the most practical one) involves levels, so fader rides should probably precede any other type of automation.

Here are just a few things we can automate. Some of them have been mentioned in this book already, while others can be heard on some commercial mixes:

- Raise the level of a specific instrument during the chorus, then bring it down again during the verse. Likely candidates are vocals, kicks, snares, and guitars.
- Mute some instruments early in the song, then introduce them later.
- Make an instrument brighter or darker during specific sections.
- Ride the level of overheads during crash hits.
- Ride the level of overheads or any other instrument up and down with relation to the tempo.
- Pan something to one place during one section, then pan it somewhere else during another.
- Apply some interesting vocal effect momentarily.
- Change the timbre of the kick during some sections.
- Mute the double-tracked vocals at points.
- Change the reverb of the snare between sections or just alternate it between hits.
- Bring down the level of some instruments to clear some space for another instrument.
- Introduce distortion on the bass during the chorus only.
- Increase the compression on the drums as the song progresses.
- Widen and narrow the stereo width of an instrument during various sections.

Automation engines

Automation engines work on the principle of storing the position of controls using automation events. An automation event typically includes which control has been automated and its position at specific timestamps (very similar to the way MIDI systems handle MIDI control messages). Before any automation has been written, each control is free to move. Once even a single automation pass has been performed, the control position is often bound to the calculated position between two automation events or the position of the latest automation event.

It is worth knowing that some automation systems store events with SMPTE time code, which advocates the use of 30 fps for higher-resolution automation. This is subject to the project not involving any video that might dictate a different frame rate.

The automation process

Performing vs. drawing

When *performing* automation, we move controls during playback. Audio sequencers (and some digital recorders) provide a graphical display of automation events and also let us draw them on screen. Sometimes, drawing automation is quicker. For example, if we want to mute a specific instrument for a minute, it should take less time to draw two mute events than performing mute automation and having to wait for a minute. Drawing automation can also be beneficial when we want events to be quantized to the tempo grid. However, some argue that performing automation yields more musical results as there is always an interaction between what we hear and what we do—we respond to the music rather than approximating the effect of drawn events (which effect we mostly hear after drawing). Depending on the situation, different methods may be more appropriate, but it is worth remembering that, between performing and drawing, the former is more likely to involve feel.
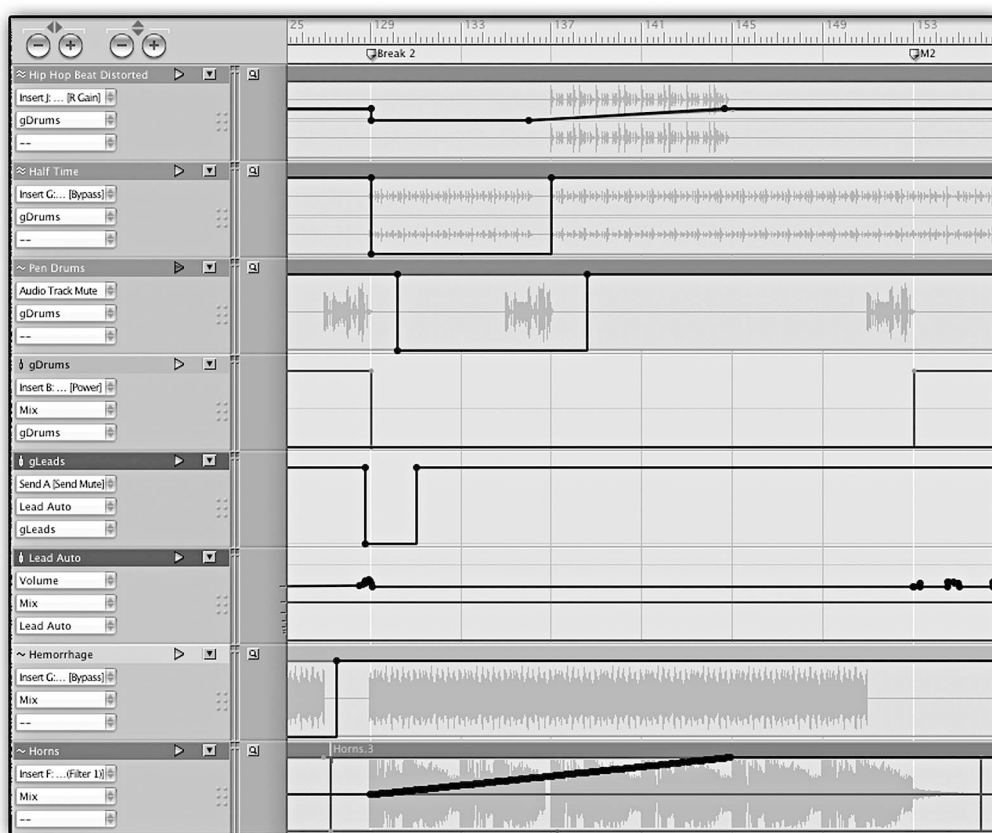
Figure 28.1 This Digital Performer screenshot shows all the automation events happening during the second break of "The Hustle." Some of these events have been drawn (such as those in the top track), and others performed (such as those in the Lead Auto track).

Performing automation

Automation is said to be written rather than recorded. We do not have to press a record button in order to write automation, although sometimes we have to tell the system which control is to be automated and sometimes we have to assign (arm) a specific channel of interest to the automation engine (more common on digital desks).

When writing automation on an audio sequencer, we either use a control surface or the mouse to alter the position of controls during playback. Ideally, automation systems would like to know when a control is touched and when it is released—often automation is only written between these two events. Some control surfaces feature *touch-sensitive controls*; these are either faders or (less commonly) rotary knobs that, by way of varying capacitance, detect a finger touch. If a control is not touch-sensitive, the automation engine starts writing automation either with the first control movement or as soon as the control position matches the existing automation value. In most cases, such control is considered released after a certain period has passed with no position changes (a period often called *touch timeout*). Controls on a computer screen are regarded as touch-sensitive—a control is considered touched as soon as the mouse button is pressed, and released as soon as the mouse button is released.

Automation modes

Automation engines may vary in their modes, features, and response to user action. We can, however, generalize about a few automation modes typical in many systems. The modes in this list are illustrated in Figure 28.2.

• **Off**—automation data is neither read nor written.
• **Read**—previously written automation is read, but new control changes are not written.
• **Touch**—new automation data is written as long as a control is touched, otherwise previous automation is read.
• **Latch**—automation is written from the moment a control is touched until playback stops.
• **Write**—new automation information is written, as long as the playback is running, overriding previously recorded automation. To prevent unwanted automation over-rides, automation engines often switch to a different mode after each pass in write mode.

Apart from when the playback stops, the writing of automation might also stop if the mode is set back to read or if the specific channel automation assignment is disarmed. Two things can happen when automation writing stops during playback: either the control jumps instantly to the previous automation position (which could generate clicks) or it slides to that position. Often the time it takes a control to slide between the two positions is called the *match period*.

Another mode known as **Trim** or *Relative* mode usually applies to levels only (faders and sends). It is useful when we want to adjust the level of previously written
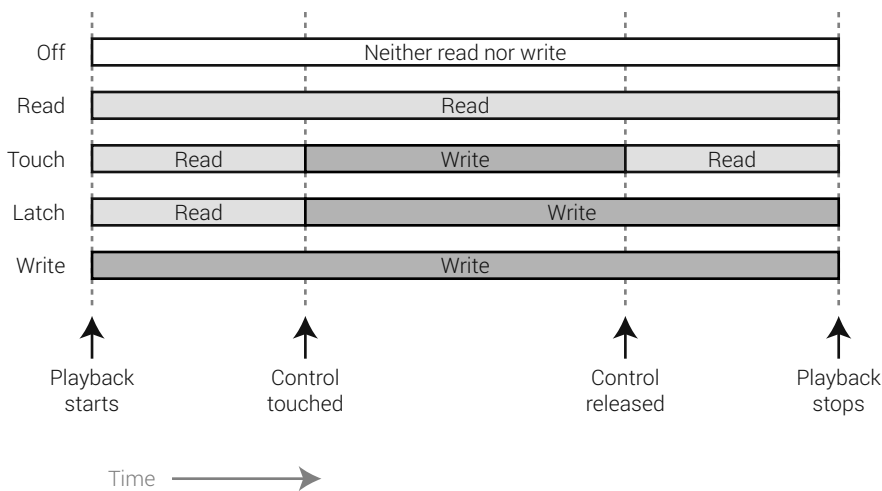
Figure 28.2 The five typical automation modes.

automation. The idea is that, instead of writing the absolute level values, trim mode simply offsets the existing automation levels by the number of dB we move the fader by. For example, if the fader is brought down from –6 to –12 dB, all automation events during the writing pass would drop by 6 dB. Systems often provide the functionality to apply the relative change throughout the song.

Automation is often one of the last things we do in a mix. After recording automation for a specific control, any adjustments can be something of an effort since we have to adjust the full automation information rather than just move the control position. This can be especially annoying if we want to alter the level of tracks after writing level automation. This is exactly what trim mode came to solve, but even trim mode can be cumbersome at times. An elegant solution to this problem was mentioned earlier—we could insert a gain plugin and use it to perform any level automation. The track's fader would then be free for global-level adjustments. Another solution is to send different instruments to an audio group and automate the level of the group instead of that of the original tracks.

Automation alternatives

Duplicates

In some situations, we might want to apply more than a few changes to a specific instrument. For example, during the chorus we might want to make the overheads louder, compress them more, narrow their stereo width, and alter their equalization. We can automate all the related controls, but there is a quicker method: most audio sequencers allow more audio tracks than any project requires. In scenarios where serious changes are needed, it sometimes pays to duplicate the track in question, trim or mute the respective sections on the two tracks (for example, having the choruses on the duplicate only, and removing them from the original track), and mix the duplicate differently, with all the involved changes (Figure 28.3).
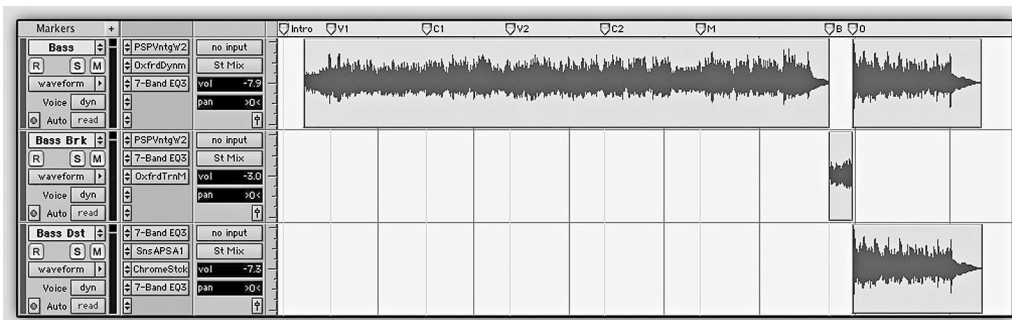
Figure 28.3 Duplicates "automation." This screenshot shows three bass tracks. The top track plays throughout most of the song and provides the main bass sound in the mix. During the break, a different sound was sought, involving a variation of tonal characteristic and level, and some additional processing. To achieve this, the bass was moved onto a new track with different processing and level (Bass Brk) for the length of the break only. The bottom track is a distorted layer that is only mixed with the main bass track during the outro.

Fades

Mix-fades are another often-overlooked practice—instruments very often fade in or out rather than instantly starting or stopping. The risk of clicks exists with any mute automation or region-trimming, whereas fades are click-proof. Transitions between sections are gelled using fades, whereas both mute automation and region-trimming can come across as very unnatural. While crossfades are more of an editing affair, fade-ins and fade-outs are a powerful mixing tool. We can achieve fades using level automation, but automation makes level adjustment a longer procedure, and where fades are needed is often realized very early in the mixing process. The fade tools provided by every audio sequencer are used for this task.

Control surfaces

When Digidesign launched the *Icon D-Control* (Figure 28.4) back in 2004, eyebrows were raised as to whether there was a place for a pure control surface as big as a large-format console. But for many engineers who used large-format analog consoles, it was justified.

The analog versus digital debate goes beyond audio quality to the realm of human interaction. Analog consoles provide the highest level of accessibility—all the controls are laid out in front of you and any sonic action you'd like to perform is within reach. In most cases, 2 or 3 seconds is all it takes to translate your sonic vision into sound. Then there are the layer-based digital desks, where usually a third or half of the mix might be readily accessible, although you might have to press the select button in order to equalize a specific track or navigate some on-screen menus in order to tweak a certain effect. On the bottom of the accessibility ladder comes a computer system, where the whole mix has to be channeled through a rather primitive device called a mouse.

Figure 28.4 The Digidesign *Icon D-Control*. This product, which looks like a large-format console, is a pure control surface for Pro Tools. Products of this kind enable a mixing experience that was once reserved for users of large-format consoles, with all its technical and creative benefits.

Source: Courtesy of Digidesign. Photo: Bill Schwob.

Compared to large-format desks, even a 21-inch screen is tiny. The more tracks a project involves, the slower the mixer navigation becomes. Plugin windows have to be opened and closed. Things can take time. Had it only been a matter of more or less time, perhaps control surfaces as large as the *Icon* would not exist. But the creative flow can easily be restrained while the brain is busy operating the computer. Audio sequencers provide a few features that let us shorten the time between our vision and its implementation. But there is no argument about the fact that mouse-mixing will never be as fast as mixing on a large-format analog console or on large control surfaces such as the *Icon*.

It goes beyond that, even. Most people find sliding faders and turning knobs much more natural than dragging a mouse on a screen, especially when it comes to automation. Then there is also the fact that a mouse is a serial interface—rarely can we change more than one control at a time, although in mixing it is frequently necessary. We might, for example, alter the ratio and threshold of a compressor simultaneously, we might fancy boosting the highs on one channel while attenuating them on another,

Figure 28.5 The Frontier Design Group *AlphaTrack*. Despite its compact size, a control surface such as this gives automation in particular and mixing in general a far more natural feel.

or we might want to bring one fader down while bringing another up during an automation pass. There are many examples.

Control surfaces come in various forms and sizes: from the large *Icon*, to the moderate-size designs such as the eight-fader Euphonix *MC Mix*, to compact designs such as the *AlphaTrack* in Figure 28.5. Many of them can be cascaded, and the larger the work surface we are using, the more accessible our mix becomes and the faster we can realize our sonic vision. But even compact surfaces such as the *AlphaTrack* can make automation and other aspects of mixing a far more natural experience.

Part III

Sample mixers

Regarding the naming conventions of tracks in this part:

Many of the audio tracks in this part involve both soloed and non-soloed versions of the instrument and treatment in question. Audio tracks with "(i)" denote the soloed (isolated) version, while "(m)" denotes the non-soloed (mix) version.

29 "Show Me" (rock 'n' roll)

Performed by The Raid.
Mark Morris: drums.
Paul Taylor: lead guitar.
Sam Street: bass.
Andy Maguire: rhythm guitar.
Adam Robinson: vocals.

Produced by The Raid and Jacques De Kooning.
Recorded at Orgone Studios, London.
Engineered by Jaime Gomez Arellano.
Edited by Ori Kadishay and Andrew Phillis.
Mixed (in Pro Tools) by Roey Izhaki.
Mastered by Jaime Gomez Arellano (www.orgonestudios.com) at Orgone Studios, London.

The Raid is a five-piece rock 'n' roll band from Hemel Hempstead, United Kingdom. An old number of theirs—"Show Me"—was rerecorded for this book at Orgone Studios in London over a period of approximately 6 hours.

"Show Me" called for a mix associated with the words *loud* and *aggressive.* A production that sprang to mind as a rough reference was "Search and Destroy" by Iggy and the Stooges, where a lot of sounds go "above red." Yet the general plan for "Show Me" was to create a more contemporary sound and not to go to similar extremes.

One specific aim for the particular mix presented here was to have it done "the analog way." This is less in the sense of how it sounds, but more in the sense of how many and which processors and effects are used. I've tried to limit these to what one might find in a typical analog studio—a dynamic section and four-band EQ per channel, and a few reverb units. The line-trims allowing the generation of distortion on an analog console were simulated by various distortion plugins. Also, to keep the mix "real," no drum triggers were used, nor was there any re-amping of the DIs. My only sin in that respect was the usage of some tom-runs.

Track 29.1 is the mix-ready of the multitrack, involving rough levels and panning only. Track 29.2 is the mixed version, and Track 29.3 is the mastered version.

Drums

Overheads

Inserts: Digirack *Expander/Gate*, Bomb Factory *76*, Digirack *EQ III*.

The drums were recorded with two sets of overheads: a pair of *U87*s above the drum kit in Blumlein configuration (Track 29.4) and a pair of Cole *4038*s as a spaced pair about 4 m from the drum kit (Track 29.5); there was also a single NT5 in the corridor leading to the studio (Track 29.6). The main studio, where the drums were recorded, is evidently quite big and reverberant, and virtually all the overheads had too much room for a mix that called for a rather tight space. From all the options, the *U87*s were picked for having least space and for their tonality; but these still needed an expander to reduce the amount of ambiance, while also adding some punch (Tracks 29.7 and 29.8).

As the production starts with a drumbeat only, it made sense to make it sound exciting, a task for which a pumping compression or distortion usually work when applied on overheads or a drum mix. The Bomb Factory *76*, set to fairly drastic settings (fastest release and attack and limiting ratio), did a pretty good job of adding both pumping compression and distortion, and gave listeners a reason to get excited (Track 29.9). It is worth mentioning that, due to the nature of the intro beat, this is where this production started being mixed; so, in a way, everything that followed was somewhat referenced to this intro beat.

One problem with these distorted-pumping overheads was that, while they were more than appropriate in the intro, as the song progressed, the effect made the drum sound spin out of control. So these overheads were automated to dive 3 dB when the bass ws introduced and another 7.6 dB when the guitars explode; they do come up again in the first part of the second verse to compensate for the sparser mix.

Following the compressor, an EQ was inserted, mainly in order to reduce some low-end boominess and add some clarity (Track 29.10).

Overheads (duplicate)

Inserts: Digirack *Expander/Gate*.

As the pumping overheads had to dive in level, it was necessary to introduce another set of overheads of a more relaxed and solid nature. Effectively, the previous track was duplicated, but the compressor and EQ were removed from it (Track 29.11). This duplicate track was consistent in level, and together with the previous overheads it constitutes a parallel-compression/parallel-distortion arrangement. Tracks 29.12 and 29.13 present the intro with and without the duplicate.

Kick

Inserts: Digirack *Expander/Gate*, Digitrack *Compressor/Limiter*, SoundHack *+Chebyshev*, Digirack *EQ III*.

The original kick (Track 29.14) was first gated in order to eliminate spill and to add a bit of "click" (Track 29.15); the gate was automated to be bypassed during crescendo, where it gates the soft hits. A compressor follows to add more attack (Track 29.16).

Although both the expander and compressor add click by shaping the dynamic envelope of the kick, after mastering there is a chance these will be greatly softened. So the *+Chebyshev* follows to add a click through frequencies, which is more likely to be retained after mastering. The fifth harmonic on the *+Chebyshev* was set all the way up, and, although often too extreme, in this case it "burns" the attack and gives it a somewhat analog feel (Track 29.17).

Finally, an EQ shaped the final tone of the kick, with a boost of 6.6 dB at 81 Hz to add oomph, –5 dB at 377 Hz to remove plastic, and a 3.4-dB boost at 7.3 kHz to fine-tune the attack (Track 29.18).

Snare

Inserts: McDSP *G Dynamics*, Soundtoys *Decapitator*, Digirack *EQ III*.

Sends: 112 dB *redline reverb*.

The original snare (Track 29.19) first goes through the McDSP *G Dynamics*, where the compressors have two roles: to balance out the hits and to add some attack (Track 29.20). Originally, the gate was also engaged to remove spill, but later in the mixing process it was determined that the snare sounds all the more healthy without the gate.

Following the compressors, the *Decapitator* was employed to add a few harmonics while also shaping what was a slightly wimpy tonality (Track 29.21). This was followed by an EQ with –2.9 dB at 1.2 kHz that eases a frequency range that, despite making the snare stick out a bit more, was annoying (Track 29.22).

All the reverbs used in this mix were the 112 *redline reverb* and all were loaded as "stubs"—temporary plugins that were there to give an idea but their effect was not fine-tuned. Somehow, though, these reverbs held well throughout the mix. The *redline reverb*, in my opinion at least, seems to do an excellent job of adding shine, which was very much what was needed from a reverb in a dense track such as "Show Me." In the case of the snare, the reverb was added to push the snare backwards as it was standing out a touch too much (Track 29.23). But the reverb didn't complement the drumbeat at the intro, and so the send is only enabled as the guitars explode.

Hi-hats

Inserts: Soundtoys *Decapitator*.

I rarely find hi-hat microphones useful (many times I wish hi-hats had never been invented), but this is a good example of where they can be. The hi-hats are only

mixed during the intro up until the explosion. During the first 12 bars, the hi-hat microphone (Track 29.24) was panned opposite the overheads' hi-hat to widen the stereo image and for the effect caused by the short delay between the hi-hat microphone and the overheads. Then, the hats were utterly distorted using the Soundtoys *Decapitator* (Track 29.25); this is part of the heavily processed strategy used during the intro, and there's not much harm in doing so in this section, as the distortion has little to mask.

Tom 1

Inserts: Digirack *EQ III*.

Both the (real) toms are only heard during the pre-chorus bridges. Tom 1 (Track 29.26) was only equalized with an HPF that rolls off low frequencies that added unwanted boominess; a 2.5-dB boost at 107 Hz compensates for the removed lows and –10 dB at 249 Hz improves the tom timbre (Track 29.27).

Tom 2

Inserts: Bomb Factory *76*, Digirack *EQ III*.

Tom 2 (Track 29.28) is being compressed by the Bomb Factory *76* in order to balance out the hits and also add some attack (Track 29.29). The following EQ is very similar to the treatment of Tom 1 (Track 29.30).

Toms 1 and 2 (tom runs)

Inserts: Bomb Factory *76*, Digirack *EQ III*, Soundhack *+Chebyshev*.

Both toms (Track 29.31) are treated in a very similar fashion. First, the Bomb Factory *76* adds some attack (Track 29.32). Then an EQ boosts the tone in lows and attenuates the low-mids (Track 29.33); in the case of tom 2, there is also a boost on the 2 kHz. Then the *+Chebyshev* adds some harmonic complexity and definition (Track 29.34).

Bass

Inserts: PSP *Vintage Warmer 2*, Digirack *EQ III*, Soundtoys *PhaseMistress*.

The bass (Track 29.35) first goes through the *Vintage Warmer*, which largely balances out the level of different notes.

Then the *EQ III* shapes the sound of the bass, but, more importantly, blends the bass into the mix. There is a 3.5 dB boost on 71 Hz to add some fatness (Track 29.36); a –5.6 dB dip at 194.2 Hz that removes some honkiness (Track 29.37); and a 3.4 dB boost at 2.7 kHz to add some metal and definition (Track 29.38). Although the settings ended up with what seems to be a rather typical treatment for a bass guitar, it did take some effort to reach the final frequency, gain, and *Q* settings.

Following the EQ, the *PhaseMistress* added a touch of fluidity and life to the bass (Track 29.39). Also, the bass was automated in level to drop with the intro explosion, and then was brought up again for the first part of the second verse.

Rhythm guitar

Inserts: Soundtoys *Decapitator*, Skonalsis *SV-719*.

There are two mono overdubs of the rhythm guitar and these were panned hard-left and right (Track 29.40). During the bridges, both parts were "manually gated" by cutting the regions and then adjusting fades; this was done to achieve a sharper and tighter effect (Tracks 29.41 and 29.42). Both mono recordings were routed to the rhythm guitar bus, where processing took place.

Although the rhythm was powerful all right, I felt that its sound was a bit too washy and needed more dirt. So what was already distorted was distorted again using the *Decapitator* (Track 29.43). The following *SV-719* was in ducker mode with the vocal ducking the rhythm guitar. The idea was to clear some space for the vocal but also add some variations on the otherwise flat level of the rhythm. After setting it all up, the ducking of the rhythm seemed to hamper the mix, but if compensated the rhythm seemed to be too loud when there were no vocals. So the final settings involve a tiny 1 dB of ducking, which, admittedly, was very subtle (Track 29.44).

Lead guitar

Normal lead

Insert: Digirack *EQ III*.

Sends: 112 dB *redline reverb*.

Two lead guitar overdubs were panned roughly halfway to the extremes (Track 29.54). Both are equalized very similarly with an HPF (around 280 Hz) to remove rumble (Track 29.55), a 2.5-dB boost around 450 Hz to add some "meat" (Track 29.56), and a 3.4-dB dip at 3.7 kHz to reduce masking and excess of edge (Track 29.57).

The dry lead felt a bit frontal so it was sent to the ambiance reverb, which panned it back a bit (Track 29.58).

Lead slides

Insert: Digirack *EQ III*.

During the bridges, the lead parts were copied to new tracks, where both overdubs (Track 29.59) were treated with only an HPF at 261 Hz (Track 29.60). In these sections, the lead was not sent to a reverb. The idea was to create some interest by briefly modifying the tonality and spatial location of the lead.

Lead solo

Insert: Soundtoys *Echoboy*.

Sends: 112 dB *redline reverb*, Soundtoys *Echoboy*.

Although the solo was indeed powerful (Track 29.61), I was after a slightly more aggressive type of sound. The delay part of the *Echoboy* was disabled (0 delay time), and the extra aggressiveness was the result of a high input and saturation setting going into the "master tape" simulation, which was part of the plugin (Track 29.62).

During the solo, the lead was sent to the same reverb as the vocal; this was just an experiment, and I liked the result (Track 29.63). Then, during the second part of the solo, the lead was also sent to another *Echoboy*, but this time one that was configured as a delay. This was done to add a bit more chaos and to fill some space shortly after the solo ends (Track 29.64).

Lead accents

Insert: Soundtoys *Echoboy*.

Sends: 112 dB *redline reverb*.

In three instances, I wanted to emphasize the part played on only one of the lead overdubs, so the selected part was moved to a new track and treated with the same effect as the solo guitar while being panned center (Track 29.65). The lead in these parts was sent to the ambiance reverb for depth purposes (Track 29.66).

Feedbacks

Sends: 112 dB *redline reverb*.

There are two feedback tracks during the intro and three during the second verse (Track 29.67). All of these had a generous send to the ambiance reverb (Track 29.68).

Vocal

Having listened to the song before, it seemed obvious to me that it demanded a different vocal sound in the second verse. I also thought the voice on the choruses should be more aggressive than that of the verses. Perhaps having yet another effect on the vocal in the bridges was overdoing it slightly, while for some it might seem nothing of the sort.

So the vocal ended up having four different effects. Each was achieved on a different track: there's the effect only heard on verse one; there's the effect heard on the first and third bridges; there's the effect on the choruses; and there's an effect on verse two, which carries on to bridge two (so as not to have all three bridges containing the same effect).

Verse one vocal

Inserts: Sonnox *Oxford Dynamics*, Soundtoys *Decapitator*, Soundhack *+Chebyshev*.

Sends: 112 dB *redline reverb*.

The vocal on verse one (Track 29.69) was first balanced out using the *Oxford Dynamics*, in which the side-chain was equalized to reduce low-frequency content (Track 29.70). It was then distorted using the *Decapitator* (Track 29.71), and then the *+Chebyshev* added additional dirt and harmonics (Track 29.72). Finally, the vocal was sent to the vocal reverb, which placed it in space (Track 29.73).

Bridge vocal

Inserts: Sonnox *Oxford Dynamics*, Soundtoys *Echoboy*, Digirack *EQ III*.

Sends: 112 dB *redline reverb*.

The same *Oxford Dynamics* compression was applied on the bridge vocal (Track 29.74). This was followed by the *Echoboy* with its delay set to 0, but the saturation module set to high input and saturation into a telephone simulation (Track 29.75). The *EQ III* attenuated a 4-dB excess of 2 kHz (Track 29.76). There was more reverb on this vocal, which sent it backwards and gave it an announcer-type effect (Track 29.77).

Chorus vocal

Inserts: Sonnox *Oxford Dynamics*, Soundtoys *Decapitator*, Digirack *EQ III*.

Sends: 112 dB *redline reverb*.

The same compressed version of the vocal (Track 29.78) fed the *Decapitator* in high-drive mode and resulted in a very obvious distortion, suitable for this part of the song (Track 29.79). Then the EQ attenuated a 4-dB excess at 2.6 kHz (Track 29.80). There was then some reverb (Track 29.81).

Verse two vocal

Inserts: Sonnox *Oxford Dynamics*, Soundtoys *Decapitator*, Digirack *EQ III*.

Sends: 112 dB *redline reverb*.

The same compressed vocal (Track 29.82) fed the *Decapitator*, but this time I was after a crisp effect (Track 29.83). The EQ was added in order to give it a unique resonance. A 14-dB boost at 852 Hz was balanced by a –10-dB attenuation at 1 kHz. As the two bands overlap, the boost on 852 Hz was pulled down to 7 dB (and the –10 dB was pulled up to 3 dB); the effect can be heard on Track 29.84, and with reverb on Track 29.85.

30 "It's Temps Pt. II" (hip-hop/urban/ grime)

"Ain't hip-hop, ain't urban, ain't grime … it's temps …," say the lyrics. Indeed, "It's Temps Pt. II" does not fall directly into these genres, but it involves a fresh approach that has its roots in all three. Without a doubt, the most distinguishing element in this production is the complex beat. I have taken the liberty of mixing this distinctive beat in a distinctive way, trying to craft a less-than-usual sound. As is typical of many hip-hop mixes, the beat and the vocals are the life and soul of the production, with all other elements tailored behind them.

Track 30.1 is the producer's rough mix. Track 30.2 is a mix-ready version involving rough levels and panning, but no processing or effects. Track 30.3 is a snapshot of the mix-in-progress. Track 30.4 is the final mix, and Track 30.5 is the mastered mix.

Beat

Main kick

Inserts: UAD *Cambridge EQ*, Cubase *Compressor*.

The main kick (Track 30.6) first goes through the *Cambridge EQ*, where a single parametric filter (310 Hz, –6.8 dB, *Q* 1.05) attenuates low-mids that contributed little to the overall kick sound—with less mids, the kick appears to have more oomph and more attack (Track 30.7). The EQ is followed by the Cubase *Compressor* [threshold –20 dB, ratio 2:1, peak sensing, attack 67.4 ms, release 10 ms], which adds some punch (Track 30.8).

DOI: 10.4324/9781003303077-34

Second and bat kick

Neither the second kick (Tracks 30.9–30.11) nor the bat kick (Tracks 30.12–30.14) were treated. They are both simply layered underneath the main kick.

Main snare

Inserts: Sonnox *Oxford Dynamics*, t.c. electronic *EQSat*, SoundHack *+Decimate*.

Sends: Snare Reverb (Cubase *RoomWorks*).

The main snare (Track 30.15) first goes through the *Oxford Dynamics*, where a compressor [threshold –16.8 dB, ratio 2:1, hard-knee, attack 2.78 ms, release 127 ms] adds some attack (Track 30.16). The compressor is followed by the *EQSat*, where three bands are employed: a parametric filter [126.7 Hz, 3.5 dB, bandwidth 0.63 oct] adds some body (Track 30.17), another parametric filter [902.1 Hz, 4.5 dB, bandwidth 1.6 oct] adds some presence and attack (Track 30.18), and a high-shelving filter [1.45 kHz, 1 dB] adds some highs as part of frequency tuning (Track 30.19). Following the EQ, *+Decimate* reduces the bit depth to 8.04 bits. This adds some definition-noise and gives a touch of the famous 8-bit sample sound (Track 30.20). The level of the snare was automated to rise during the set of choruses after the second break. The main snare is also sent to the *RoomWorks* reverb, where a very short reverb sends it backward from its front position, which I found was too in-your-face (Track 30.21).

Second snare

Inserts: PSP *MasterQ*.

Sends: Snare Reverb (Cubase *RoomWorks*).

The second snare (Track 30.22) is not a layer of the main snare, but an independent instrument that plays on different beats (Track 30.23). This track only goes through the *MasterQ*, where a parametric filter [306 Hz, –8.41 dB, *Q* 0.53] pulls a wide range of mids that only muddled the sound while covering an important range for other instruments (Track 30.24). This track is also sent to the snare reverb (Track 30.25).

Claps

Inserts: Sonnox *Oxford Dynamics*, Cubase *StudioEQ*, SoundHack *+Decimate*. The original claps (Track 30.26) sounded too light and too natural. The initial burst of claps is usually sufficient for listeners. There is little damage done by shortening claps (especially in a busy mix), a practice that can also make claps snappier. The gate [threshold –8.4 dB, range –41 dB, attack 5 μs, hold 35 ms, release 9.2 ms] of the *Oxford Dynamics* was employed for this task (Track 30.27). It is followed by the *StudioEQ*, where three bands are used: an HPF [357 Hz] rolls off dispensable low-frequency content (Track 30.28), a parametric filter [2,057 Hz, –3.1 dB, *Q* 0.5] further reduces high-mids that contribute little to the mix (Track 30.29), and a high-shelving filter enhances definition and tunes the claps into the frequency spectrum (Track 30.30). The EQ is followed by the *+Decimate* [8.04 bits], which, similar to its function on the snare, adds some grit and 8-bit sound (Track 30.31).

Hi-hats

The hi-hats track (Track 30.32), which plays through most of the production, was left untreated. The track already contained some spatial information that sent it backward in the depth field. Being an additional element to the beat, it was panned left to around 10:00.

Wood

Another beat element that was left untreated is the wood (Track 30.33). It was panned around 14:00.

Broken bells

The broken bells (Track 30.34) suffer from exaggerated harshness and lack of grace. I found this track somewhat disturbing, and in my view the mix was better off without it. It was omitted.

Cajon

Intro/Break Inserts: t.c. electronic *Filteroid*, PSP *608 MultiDelay*.

Other Sections Inserts: PSP *MasterQ*.

One important percussion element in this tune is the cajon. It plays throughout most of the song, but has special importance during the intro and the break. This led me to split the source track into two, having a dedicated track for the intro and break and another for the rest of the song.

The intro/break cajon (Track 30.35) first goes through the *Filteroid*, where a resonant LPF is bar-sync modulated (Track 30.36). Then it flows through the PSP *608*. On the multitap delay, the mix control is set to 18 percent wet. Two taps are operational—one set to eighth-note, another to quarter-note. Both taps have a feedback set to around 50 percent and both have an HPF in their feedback loop (Track 30.37). Just before the first main beat drops in, the amount, gain, and balance of the first tap are automated. This creates a sustained feedback that sweeps from left to right, and could also appear to sweep up and down (Track 30.38). The level of the cajon track was automated to rise during the break.

The cajon that plays throughout the rest of the song is layered underneath the main beat. It is only treated with the *MasterQ*, where an HPF [97 Hz] clears the lows (Track 30.39).

Beat group

Inserts: PSP *Vintage Warmer*, UAD *Fairchild*, UAD *1176 SE*.

All beat tracks but the main kick are routed to an audio group. The reason for the kick's exclusion is that I felt it would benefit from independent processing—not being affected by the heavy compression taking place on the beat group. A quick look at the inserts chain reveals that the beat is processed by three compressors connected in series. Here is

what happened. The first compressor in the chain was the *Vintage Warmer*, configured with parallel compression settings [mix 50 percent, drive 0 dB, knee 100 percent, speed 100 percent, auto release]. Its main role is to condense the beat and add some punch (Tracks 30.40 and 30.41). The *Fairchild* [input gain 7 dB, threshold 3, fastest time constants] was loaded for the purpose of enhancing the compression effect and to give some unique character (Track 30.42). The *1176* was loaded as a possible alternative to the *Fairchild*, but when it was loaded with the *Fairchild* still on, the result of the two working in series was infused with character (Track 30.43). The *1176*, however, is not operational throughout the song since it made the beat somewhat slim; instead, it was automated to take effect during the last four bars of the first verse and throughout the second verse.

Break beat

Inserts: Sonnox *Oxford Dynamics*.

Sends: Break Reverb (t.c. electronic *ClassicVerb*, Sonnox *Oxford Dynamics*, PSP *Nitro*, Sonnox *Oxford Dynamics*).

The break beat (i.e., the beat during the break) called for a change of theme and more drama. I have chopped the kick, snare, and hi-hats from their main tracks and moved each to a new track. The kick was equalized with the PSP *Neon* (Tracks 30.44 and 30.45), the snare was processed with the PSP *MasterQ* (Tracks 30.46 and 30.47), and the hats were left untreated. These three tracks were then routed to an audio group, where the *Oxford Dynamics* added some compression in linear mode (Tracks 30.48 and 30.49).

But the powerful effect on the break beat is the outcome of an inserts chain that succeeds the reverb to which all three tracks were sent. The reverb is based on a modified "Small Hall" preset from the *ClassicVerb* (Track 30.50). The reverb is followed by the *Oxford Dynamics*, where the gate gates the reverb and creates chaotic level variations (Track 30.51). The combination of bit reduction and a wave shaping from the "PumpItUp" preset of the *Nitro* adds a generous amount of distortion (Track 30.52). Then another *Oxford Dynamics* compresses the sound, using the more hectic linear mode (Track 30.53). Finally, the second plugin in the chain—the gate—is bypassed shortly before the end of the break, to leave an audible reverb decay just before the chorus kicks in (Tracks 30.54 and 30.55).

Bass

Main bass

Inserts: Noveltech *Character*, PSP *Neon*, Sonnox *Transient Modulator*.

A single bass track was submitted with the multitrack. The original track (Track 30.56) had a generous amount of mids and highs that were meant to increase the bass definition, but I found these specific frequency ranges unappealing, so I decided to filter them. The first plugin the bass goes through is the Noveltech *Character*. The enhancer was inserted to fatten the lows of the bass, although as a consequence the highs were

also enhanced (Track 30.57). Next comes the *Neon*, where in linear-phase mode three bands are used. First to be applied is an LPF [380 Hz, 12 dB/oct], which filters the mids and highs of the bass. This makes the bass a pure bass, without the mids and highs (Track 30.58). A low-shelving filter [100 Hz, –1.06 dB] that controls the subs was added later in the mix as part of frequency tuning (Track 30.59). Also, a parametric filter with a narrow *Q* [99 Hz, –5.18 dB, *Q* 11.25] was employed to shape the sound of the bass by removing some tone from around 100 Hz (Track 30.60). Following the EQ, the *Transient Modulator* adds some attack to the bass—this accents the initial level burst and lets the bass stand out better in the mix (Track 30.61).

Bass support

Virtual Instrument: Access *Virus TI*.

Filtering the mids and highs from the bass removed some problematic frequencies, but then the bass lacked definition. So an email was sent to the producer, who replied with the MIDI bass track file. The file was loaded onto an instrument track, with the Access *Virus TI* as a virtual instrument. The level of this support track was automated to have various degrees of effect at different sections. For example, it was attenuated in level during the intro and second verse, where the sparser arrangement made the added distortion highly noticeable.

Two slightly modified *Virus* presets were used in this arrangement. "Dukbass+BC" is layered with the bass throughout the production (Tracks 30.62 and 30.63). Rob Papen's "Vacin RP" is essentially a distorted lead, wide in image, that only plays during the intro and the choruses. Its addition made the bass bigger and more defined during these important sections (Tracks 30.64 and 30.65).

Other tracks

Steel drums

Inserts: PSP *MasterQ*.

Final Inserts: PSP *MasterQ*, Cubase *MonoToStereo*, PSP *Nitro*.

Steel drums play a key role during the intro and break sections, during which the track level was automated to rise. Two versions were submitted with the multitrack—one with and one without delay (Tracks 30.66 and 30.67). I chose to use the one with delay. The original image of the raw stereo track was too wide for the mix and I wanted the instrument to be panned left. To achieve this, the image was narrowed by moving the right pan-pot to the center (Track 30.68).

The track was equalized by the *MasterQ*, where three bands are used: an HPF [367 Hz, 12 dB/oct] rolls off dispensable lows (Track 30.69), a parametric filter [2.3 kHz, –6.6 dB, *Q* 0.26] results in a more rounded sound (Track 30.70), and a shelving filter [9 kHz, –5 dB] tunes this track to the frequency spectrum and helps send it slightly backward (Track 30.71).

The steel drums also conclude the production, and I wanted the effect during the final seconds to be more unique than in the raw track. So on a duplicate track, which only plays once the music drops out, the same *MasterQ* was inserted with the same settings as above; it was then followed by the Cubase *MonoToStereo*, which actually sums the stereo track to mono (Tracks 30.72 and 30.73); then it is followed by the *Nitro*, which applies an autopan effect synced to the tempo of the song (Track 30.74).

Gtr Pizz

Inserts: PSP *Nitro*, PSP *MasterQ*.

Sends: Gtr Chorus (UAD *DM-1 Delay Modulator*).

Gtr Pizz (Track 30.75) was first processed by the *Nitro* and its "PumpItUp" preset—a combination of a wave shaper, bit reduction [10.22 bits], sample rate reduction [19.5 kHz], and an HPF [748 Hz]. This combination adds some harmonics that fatten the sound (Track 30.76). The *Nitro* was followed by the *MasterQ*, where an HPF [277 Hz, 12 dB/oct] rolls off an excess of lows (Track 30.77). The original mono track sounded small and panning it to either extreme could create a stereo imbalance. Thus, it was sent to the *DM-1*, where a chorus [left delay 23.3 ms, right delay 68.1 ms, modulation rate 0.67 Hz, modulation depth 10 percent, feedback –29 percent, damping 3.7 kHz, 100 percent wet] adds a noticeable stereo size (Track 30.78).

French horns

Inserts: PSP *MasterQ*, PSP *Nitro*, PSP *84*.

French horns (Track 30.79) play alongside Gtr Pizz on quite a few occasions. The track was first processed by the *MasterQ*, where an HPF [234 Hz, 24 dB/oct] rolls off an excess of lows (Track 30.80), a parametric filter [754 Hz, –3.2 dB, *Q* 2.8] lightly cleans the tonality of the instrument (Track 30.81), and another parametric filter [4.4 kHz, 3.6 dB, *Q* 2.8] adds some definition that lets the instrument stand out more in the mix (Track 30.82). Following the equalizer, the *Nitro* with the same "PumpItUp" preset adds distortion and some flesh (Track 30.83). The original stereo track included a stereo effect that did not translate well when summed to mono (Track 30.84). The PSP *84* was employed to add 41 ms of delay to the right channel only. This not only resolved the mono-summing issue, but also gave a more distinct stereo effect bound to the extremes (Tracks 30.85 and 30.86). The image of Gtr Pizz is narrower than that of French horns, so when the two play together the stereo image appears to widen. It is then narrowed when Gtr Pizz plays alone (Track 30.87).

Big stack

Inserts: PSP *MasterQ*.

Sends: Big Stack Delay (Cubase *MonoDelay*).

Big stack arrived on the multitrack as a mono track (Track 30.88). It was processed by the *MasterQ*, where a single parametric filter with a wide *Q* [6.2 kHz, 1.5 dB, *Q* 0.41] adds some definition and tunes the instrument to the frequency spectrum (Track 30.89).

Quite clearly, panning this track center would make it clash with more important tracks such as the beat and vocals. Panning it to either side could cause stereo imbalance. A stereo effect had to be added, and a delay was chosen [eighth-note delay, feedback 18 percent, feedback LPF at 15 kHz]. The original track was panned to the left around 10:00 and, although the delay could have been mirrored to 14:00, experimentation resulted in it sounding better panned farther right, to around 15:00 (Track 30.90).

To add some variation and create a better stereo balance, the original track was duplicated and the duplicate was panned to the right around 14:00, with another delay panned to 9:00. But the original and the duplicate never play together—separating each hit on both and muting each hit on either meant that either the original played with its right delay or the duplicate with its left delay. The mute arrangement was not based on an alternating left–right pattern, but on a more experimental arrangement (Track 30.91). The level of the duplicate was automated to make its perceived level more consistent. In addition, both the original and the duplicate were sent to an audio group for the sole purpose of level control and automation.

Dist Bells

Sends: Dist Bells Reverb (UAD *RealVerb Pro*).

The Dist Bells track was left untreated, but the dry track (Track 30.92) needed to be sent backward in the mix and called for a nice long reverb. It was awarded with its own reverb—the *RealVerb Pro*. The "Big Bright Hall" preset was selected and its decay time was lengthened to 14.5 seconds (Track 30.93). When heard in isolation, this long decay seems to take forever to diminish, but not so when the mix is played along. Also, the noise component of the distorted bells clearly lingers on the reverb tail when heard in isolation. Although this could have been rectified, I found this addition of sustained noise to be a sort of high-frequency halo to the mix, and decided to keep it. The level of Dist Bells was automated between various sections.

LPTwin

Inserts: Sonnox *Oxford Dynamic*, PSP *MasterQ*.

Sends: Ambiance Reverb (UAD *DreamVerb*), *LPT Delay* (PSP *Lexicon 42*).

LPTwin is essentially a synthesized guitar (Track 30.94). The original track presented me with a common problem where low notes got lost in the mix while high notes stood out. The compressor on the *Oxford Dynamics* [threshold –9.9 dB, ratio 4:1, hard-knee, attack 5.2 ms, release 127 ms] was employed to level out the various note levels (Track 30.95); this was done with help from side-chain equalization involving a drastic low-shelving filter [1 kHz, –20 dB] and a drastic parametric filter [3.2 kHz, 17 dB, *Q* 1.287]. The compressor is followed by the *MasterQ*, where an HPF [229 Hz, 12 dB/oct] rolls off lows that only muddy the mix (Track 30.96) and a parametric filter [4 kHz, 1.31 dB, *Q* 0.85] adds some definition (Track 30.97).

LPTwin is sent at low level to the ambiance reverb just to blend it into the mix ambiance (Track 30.98). The ambiance reverb is the UAD *DreamVerb* and its slightly

modified "Hall Cathedral" preset. LPTwin is also sent to the PSP *Lexicon 42*, which adds a distinguished WildWestlike effect. While LPTwin is panned around 14:30, its delay is panned nearly center. This creates a wide enough image and a fine sense of depth (Track 30.99).

Trumpet 1

Inserts: UAD *DM-1 Delay Modulator*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

On the submitted multitrack, there was only one trumpet track, but I decided to split it into two tracks and process each differently. Trumpet 1 involves the staccato sections, while Trumpet 2 provides two legato sections. Trumpet 1 (Track 30.100) is first processed using the *DM-1*, where a chorus [left delay 44 ms, right delay 44.9 ms, modulation rate 0.67 Hz, modulation depth 10 percent, feedback 20 percent, damping 7 kHz, 25 percent wet] adds some dimension and sends the instrument backward (Track 30.101). The track is also sent to the ambiance reverb for depth positioning (Track 30.102).

Trumpet 2

Inserts: Sonnox *Oxford Transient Modulator*, Sonnox *Oxford EQ*.

Sends: Ambiance Reverb (UAD *DreamVerb*), *TrumpetVerb* (UAD *Plate 140*).

Trumpet 2 (Track 30.103) first goes through the *Transient Modulator* [ratio 0.58], which was employed to add attack so each note could cut through the mix better (Track 30.104). It is followed by the *Oxford EQ*, where an HPF [188 Hz, 6 dB/oct] rolls off dispensable lows (Track 30.105). Trumpet 1 is sent to the ambiance reverb for depth positioning (Track 30.106). I wanted a nice reverb with a nice tail on the trumpet, but increasing the send level to the ambiance reverb would also position the trumpet farther into the depth field. So, while the ambiance reverb is in charge of the front–back position, the Plate 140 was employed to add a nice dense tail (Track 30.107).

Screamin B4

Inserts: PSP *Nitro*, Cubase *VSTDynamics*.

Sends: Ambiance Reverb (OAD *DreamVerb*).

Screamin B4 (Track 30.108), a Hammond emulation, plays during most choruses and during the break. It first goes through the *Nitro*, where the "Phat Flutter" preset adds some size and distortion (Track 30.109). One issue with this track was that the last notes were too loud. The compressor of the *VSTDynamics* [threshold –12.5 dB, ratio 1.8:1, attack 1 ms, release 228 ms] was employed to contain the level of these last notes (Track 30.110). This track was sent to the ambiance reverb in order to blend it into the mix space (Track 30.111). The level of this track was automated to drop in level during the break, rise after it, then rise again during the final two choruses.

Analog King

Inserts: PSP *MasterQ*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

Analog King (Track 30.112) plays in the last choruses of the outro and becomes one of the main instruments once the vocals drop out. This track was treated with the *MasterQ*, where three bands are used: an HPF [304 Hz, 12 dB/oct] rolls off dispensable lows (Track 30.113), a parametric filter [2 kHz, –5 dB, *Q* 1.69] is in command of tone shaping (Track 30.114), and a high-shelving filter [8 kHz, 4.5 dB] adds some definition and tunes the track to the frequency spectrum (Track 30.115). This track is also sent to the ambiance reverb, which blends it into the sound stage (Track 30.116).

Imposcar

Inserts: Cubase *Studio EQ.*

Imposcar only plays during the verses. A section of this track that played during the first chorus was muted due to arrangement density. Being one of the least important elements in the production, it was mixed underneath everything else. Although a mono version was also submitted, I chose to use the stereo version of this track, which was already imprinted with autopan (Track 30.117). Imposcar only goes through the *Studio EQ*, which shapes the tonality of this track to have more emphasis on the low-mids and less on the high-mids. Two bands are operational on the EQ: a parametric filter [735 Hz, 5.3 dB, *Q* 0.5] boosts the low-mids (Track 30.118) and another parametric filter [3.3 kHz, –2.9 dB, *Q* 0.5] attenuates the highs (Track 30.119).

Vocals

The submitted multitrack consisted of 13 vocal tracks. Since some would benefit from different treatment at different sections, I split certain tracks into two. The final mix involves 21 vocal tracks in this arrangement:

- Intro:
  - vIntro1
  - vIntro2
  - vIntro3
- Verse:
  - vVerseLead
  - vVerse1 (also plays during the intro)
  - vVerse2
  - vKsh
- Chorus:
  - vChorusLead
  - vChs1
  - vChs2

- – vChsAns1
  - – vChsAns2
  - – vChsAns3
  - – vChsAns4
  - – vTemps1
  - – vTemps2
- Break:
  - – vBreak1
  - – vBreak2
  - – vFX
- Outro:
  - – vOutro
  - – vOutroDelay

Having to compress 21 vocal tracks individually is not an easy task. Instead, I dialed rough settings on the lead vocal track and copied the compressor to other vocal tracks, then adjusted each compressor as needed. Nearly all the vocal tracks were compressed using the PSP *MasterComp*—it offered the right balance between character and control. On a few tracks, extra control was needed, so the Sonnox *Oxford Dynamics* was used instead. The general compression approach was to contain level fluctuations but retain healthy dynamics.

In addition to splitting some vocal tracks, specific words, sometimes phrases, were muted at points where too many vocals obscured one another. One of the main challenges in mixing these 21 vocal tracks was getting the right panning and depth relationship between the different voices.

Vocals reverb

Inserts: t.c. electronic *ClassicVerb*.

Different vocal tracks needed to be sent to a reverb for subtle depth positioning. The reverb had to be rather transparent and also suit the vocals as a sheer effect. The ambiance reverb was not fit for this task, so a dedicated reverb aux was allocated for the vocals. The *ClassicVerb* was chosen and its "Small Hall" preset highly modified. Tracks 30.120 and 30.121 demonstrate the function of this reverb.

vIntro1

Inserts: Sonnox *Oxford Dynamics*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vIntro1 (Track 30.122) was compressed using the *Oxford Dynamics* [threshold –30.3 dB, ratio 1.55:1, soft-knee 40 dB span, attack 0.52 ms, release 5 ms] with the intention of containing level fluctuations (Track 30.123). The track was also sent to the vocal reverb for depth positioning (Track 30.124). The track was panned around 11:00, and its level was automated to drop once the bass is introduced (the point at which vIntro2 becomes the main voice).

vIntro2

Inserts: Sonnox *Oxford Dynamics*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vIntro2 (Track 30.125) was also compressed using the *Oxford Dynamics* with exactly the same settings as vIntro1 (Track 30.126). The track was also sent to the vocals reverb (Track 30.127). vIntro2 was panned center, and its level was automated to drop during the very last sentence of the intro.

vIntro3

Inserts: PSP *MasterComp*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vIntro3 (Track 30.128) is essentially chopped from vVerse2. It was compressed using the *MasterComp* [threshold –24 dB, ratio 3.36:1, hard-knee, peak sensing, attack 0.78 ms, release 100 ms] (Track 30.129). It was also sent to the vocal reverb (Track 30.130) and is panned slightly to the left, around 11:00.

vVerseLead

Inserts: PSP *MasterComp*, Sonnox *Oxford EQ*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vVerseLead (Track 30.131) was compressed using the *MasterComp* [threshold –14.8 dB, ratio 4:1, soft-knee, RMS sensing, attack 0.79 ms, release 227 ms] (Track 30.132). After the compressor, an *Oxford EQ* HPF [141 Hz, 6 dB/oct] rolls off some dispensable lows (Track 30.133). This track was also sent to the vocal reverb (Track 30.134). vVerseLead was panned center and its level was automated to rise during the very first words of the first chorus.

vVerse1

Inserts: PSP *MasterComp*, Sonnox *Oxford EQ*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vVerse1 (Track 30.135) also constitutes the fourth voice during the late intro moments. Together with vVerse2, it provides the accompanying voice to vVerseLead. It was compressed using the *MasterComp* [threshold –27.7 dB, ratio 2.8:1, soft-knee, RMS sensing, attack 0.01 ms, release 100 ms] (Track 30.136). Following the compressor, the *Oxford EQ* was inserted, where three bands are operational: an HPF [229 Hz, 6 dB/oct] rolls off some lows (Track 30.137), a parametric filter [775 Hz, 1.8 dB, *Q* 2.83] was employed as part of tonality shaping and frequency tuning (Track 30.138), and another parametric filter [3.2 kHz, –1.65 dB, *Q* 2.83] helps enhance the back panning of the voice (Track 30.139). Altogether, the equalization of vVerse1 also helps to distinguish it from vVerseLead. This specific voice was panned to the right around 14:00. It was also sent to the vocal reverb (Track 30.140).

vVerse2

Inserts: PSP *MasterComp*, Sonnox *Oxford EQ*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vVerse2 (Track 30.141) was compressed using the *MasterComp* [threshold –14.3 dB, ratio 8:1, hard-knee, peak sensing, attack 0.015 ms, release 457 ms] (Track 30.142). It was equalized using the *Oxford EQ* similar to vVerse1, only without the 775 Hz boost (Track 30.143). It was panned to the left around 10:30 and its level was automated as an additional measure to balance its level. It was also sent to the vocal reverb (Track 30.144).

vKsh

Inserts: PSP *MasterComp*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vKsh was not submitted on the multitrack. At a certain point during mixing, I noticed a percussive "Ksh" in the second verse. It turned out that this was the tail of the word "splash" from vVerse2 (Track 30.145). I then copied this "Ksh" and pasted it into a new track on the downbeat of every other bar, for the duration of the second verse only. On this new track—vKsh—a *MasterComp* with the same settings as vVerse2 compresses the vocal. To send it farther into the depth field, this track was sent at high level to the vocal reverb (Track 30.146).

The final verse vocal arrangement can be heard in Track 30.147.

vChorusLead

Inserts: Sonnox *Oxford Dynamics*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

vChorusLead (Track 30.148) was compressed using the *Oxford Dynamics* [threshold –7.5, ratio 50:1, soft-knee 40 dB span, attack 0.52 ms, release 10 ms] (Track 30.149). It was also sent to the vocal reverb (Track 30.150). It was panned center.

vChs1

Inserts: PSP *MasterComp*, Sonnox *Oxford EQ*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

Both vChs1 and vChs2 are double-tracked versions of vChorusLead. vChs1 (Track 30.151) was compressed using the *MasterComp* [threshold –14 dB, ratio 8:1, soft-knee, RMS sensing, attack 0.010 ms, release 457 ms] (Track 30.152). The compressor is followed by the *Oxford EQ*, where three bands are employed: an HPF [229 Hz, 6 dB/oct] rolls off dispensable lows (Track 30.153), a parametric filter [2.7 kHz, –2.56 dB, *Q* 2.83] helps pan the voice slightly farther away (Track 30.154), and another parametric filter

[9 A kHz, –1.5 dB] reduces some sibilance (Track 30.155). As with the additional verse voices, the EQ helps to distinguish the second voice from the lead. This track was sent to the vocal reverb, with the intention of placing it slightly farther behind the lead chorus voice (Track 30.156). It is panned to around 11:00.

vChs2

Inserts: PSP *MasterComp*, Sonnox *Oxford EQ*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

The vChs2 track (Track 30.157) was processed very similarly to vChs1. The settings on the compressor are all the same apart from the threshold, which is set to –12.26 dB, and the release, which is set to "Auto" (Track 30.158). The settings on the successive *Oxford EQ* vary only slightly: an HPF [229 Hz, 6 dB/oct], a parametric filter [4.6 kHz, –2.15 dB, *Q* 1.67], and another parametric filter [10.2 kHz, –1.38 dB, *Q* 2.83] (Track 30.159). This track is panned to 13:00—mirroring vChs1. It is also sent to the vocal reverb (Track 30.160).

vAns1 and vAns2

Inserts: PSP *MasterComp*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

Both vAns1 and vAns2 (Track 30.161) were compressed very similarly using the *MasterComp*. The compressor settings on vAns1 are [threshold –15.63 dB, ratio 8:1, soft-knee, RMS sensing, attack 0.017 ms, release 100 ms], while the following differences were applied on vAns2 [threshold –17.25 dB, release 161 ms] (Track 30.162). Both tracks are sent to the vocal reverb (Track 30.163). vAnsl was panned to around 10:30, while vAns2 was panned to around 13:30. This panning position means that the answering vocals are not too far from the center, yet their image is wider than the lead chorus vocal and its doubles. This creates a centrally panned question and an answer that appears farther to the sides.

vAns3 and vAns4

Inserts: PSP *MasterComp*.

Sends: Vox Reverb (t.c. electronic *ClassicVerb*).

Both vAns3 and vAns4 (Track 30.164) are chopped from vAns1 and vAns2, respectively. Both tracks are compressed exactly like their origin tracks. The only difference is that vAns3 and vAnd4 are panned a notch farther outward and their send to the vocal reverb is lower; thus, they appear closer (Track 30.165).

vTemps1 and vTemps2

Both vTemps1 and vTemps2 contain five instances of the word "Temps," which emphasizes the other vocals at the very end of each chorus. These two tracks, which

were neither processed nor sent to the vocal reverb, were panned around 9:00 and 15:00, respectively (Tracks 30.166 and 30.167).

The final chorus vocal arrangement can be heard in Track 30.168.

vBreak1 and vBreak2

Inserts: PSP *MasterComp*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

Both vBreak1 and vBreak2, as their name suggests, only appear during the break (Track 30.169). Both are compressed using the same setting on the *MasterComp* [threshold –8.82 dB, ratio 8:1, soft-knee, RMS sensing, attack 0.017 ms, release 457 ms] (Track 30.170). The general idea with mixing these two voices was to have them farther away in the mix, to create a contrast with the main vocals. Thus, instead of sending these to the vocal reverb, they were sent to the ambiance reverb (Track 30.171). vBreak1 and vBreak2 are panned around 10:00 and 14:00, respectively.

vFX

Inserts: PSP *MasterComp*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

This track, which on the multitrack was already distorted and imprinted with delay (Track 30.172), was compressed using the *MasterComp* in order to contain level fluctuations, and sent to the ambiance reverb to help it gel with the mix ambiance (Track 30.173).

vOutro and vOutroDelay

Inserts: PSP *608 MultiDelay*.

Based on a rough mix submitted by the producer, instances of "It's Temps" were added to the otherwise vocal-free ending of the song. This was done by cutting the final instance of this phrase from vTemps2 and pasting it into a new track called vOutro. Instead of automating a send to a delay, I duplicated vOutro and inserted it on the new track (vOutroDelay) on which the PSP *608* was inserted [100 percent wet]. All the eight taps of the delay were utilized with a quarter-note spacing between each tap. The echoes were programmed to gradually drop in level and to shift outward between left and right (Track 30.174).

Vocals group

Inserts: t.c. electronic *Dynamic EQ*, Sonnox *Oxford EQ*.

All vocal tracks, apart from vOutroDelay, were routed to an audio group. The audio group provides collective level control and processing. The first insert on this audio group is the Dynamic EQ. One issue with all vocal tracks was overemphasis in the

10 kHz area. This was made even more profound by the *MasterComp*, which was loaded on many vocal tracks. The *Dynamic EQ* was configured with a single parametric band to the frequency of 10.3 kHz, –7.2 dB of gain, and a *Q* factor of 1. The equalization was set to fully dynamic, with the threshold set to –21.7 dB, attack to 1 ms, and release to 81 ms (Tracks 30.175 and 30.176). Essentially, the *Dynamic EQ* was utilized as a de-esser in the 10 kHz area. Following the *Dynamic EQ* came the *Oxford EQ*, which was added at a later mixing stage to tune all vocals into the frequency spectrum. Two bands are used on the EQ: an HPF [67. 1 Hz, 6 dB/oct] controls the low limit of the vocals (Track 30.177) and a parametric filter [2 kHz, 2.5 dB, *Q* 1.73] adds some highs that were lost as a result of the function of the *Dynamic EQ* (Track 30.178). The vocals group level was automated in two places to balance level fluctuations.

31 "Donna Pomini" (techno)

Written and produced by TheSwine.
Mixed (in Logic) by Roey Izhaki.
Mastered by Mandy Parnell (www.mandyparnell.com) at Electric Mastering, London.

It happens, sometimes, that you get a production to mix, raise the faders to audition the multitrack, and realize that there is not much for you to do—all the tracks seem to combine pretty well and the core of the mix is already there. This lets you spend more time elevating the mix, focusing more on the general feel and nuances. Such was the case with "Donna Pomini." In all honesty, the most challenging aspect of mixing this production was getting the levels right, especially between the varying sections.

The multitrack of this nearly 8-minute tune consisted of 70 tracks, but more than a few tracks involved dry and wet versions of the same instrument. The final mix consisted of 59 audio tracks, some of which only play for a few seconds. It is worth mentioning the sections in this production, as marked by the producer: intro, lead-in, break 1, main 1, straight, ocean, break 2, main 2, outro.

Track 31.1 is the mix-ready version. The mix took so little time to complete that there is only one snapshot of the mix-in-progress, presented on Track 31.2. There are not many changes between this mix-in-progress and the final mix presented on Track 31.3, but these changes are noteworthy. Track 31.4 is the mastered mix.

Ambiance reverb

Inserts: t.c. electronic *ClassicVerb*.

"Donna Pomini" is a classic example of a mix where one reverb satisfies nearly all the mix reverb needs. There are very few other reverbs in the mix and most of them only appear occasionally. Productions of this type call for a reverb that creates a sense of space while also governing depth positioning, but the reverb itself does not need to be too evident since most of the action occurs at the front of the mix. The chosen ambiance

reverb was the t.c. electronic *ClassicVerb*. The reverb itself was based on the "Classic Hall" algorithm and configured to have 0 ms of pre-delay. Additionally, the low- and high-color parameters were set to maximum and minimum, respectively—providing extra warmth and size but little definition. The tail of this reverb can be heard on Track 31.5, while mixes with and without the reverb can be heard on Tracks 31.6 and 31.7.

Beat

Out of the 59 audio tracks this mix involves, 32 are dedicated to the beat. In addition to the main beat that plays during the main sections, there are also dedicated beat tracks for the lead-in and the straight sections.

Main kick

Inserts: Logic *Compressor*, UAD *Cambridge EQ*.

The main kick (Track 31.8) was compressed using the Logic *Compressor* [threshold –11 dB, ratio 2.4:1, peak sensing, knee 0.5, attack 10.5 ms, release 92 ms], which adds punch and power (Track 31.9). On the subsequent *Cambridge EQ*, there is one parametric filter [4.4 kHz, 2.8 dB, *Q* 1.12], which adds attack and makes the kick even more prominent in the mix—a prominence customary in such productions (Track 31.10).

The compression and the added attack on what was already a powerful kick resulted in a massive sound. But during the first few bars of the production, this made the kick appear completely detached from the sparse and relaxed arrangement. One way to combat this would be automating the level of the kick, but changing its tonality provided a more creative alternative. So two more bands on the *Cambridge* were employed that are only operational during the first few bars. On one band, a parametric filter with a wide *Q* [991 Hz, –6.8 dB, *Q* 0.52] reduces much of the kick's impact (Track 31.11); this filter is bypassed with the introduction of the hats. The second band involves an HPF [34.7 Hz, 24 dB/oct], which effectively starts rolling off frequencies below 250 Hz (Track 31.12). Its role is to attenuate much of the powerful lows and subs. It is bypassed just before the introduction of the snare. The bars involving these two EQ automation events are presented in Track 31.13. The level of the main kick was automated to rise just before the first main section.

Pillow kick

Inserts: Sonnox *Oxford EQ*.

Pillow kick only appears in the two closing bars of the intro. The raw track (Track 31.14) was treated with the *Oxford EQ*, where a single parametric filter [4.3 kHz, –4.4 dB, *Q* 2.83] attenuates some click to result in a more pillow-like sound (Track 31.15).

One-hit snare

Inserts: PSP *MasterQ*, Logic *Noise Gate*, Logic *PlatinumVerb*.

One-hit snare (Track 31.16), as the name suggests, only involves one hit, which concludes the intro section. This track was first treated with the *MasterQ*, where three

bands shape the tonality of the snare: an HPF [73 Hz, 12 dB/oct] removes dispensable lows (Track 31.17), a parametric filter [470 Hz, –3.43 dB, *Q* 0.26] attenuates some body (Track 31.18), and a high-shelving filter [10 kHz, 2 dB] adds a touch of definition (Track 31.19). Following the EQ comes the *Noise Gate* [threshold –17 dB, –100 dB reduction, attack 3 ms, hold 130 ms, release 10 ms, hysteresis –3 dB], which shortens the length of the snare so it fits better into the rhythm (Track 31.20). This gate was configured with the next *PlatinumVerb* already in place. On the reverb, which was added as an insert rather than a send, the dry control is set to 100 percent and the wet to 45 percent. The reverb adds an audible effect and its tails smoothes the transition to the lead-in section (Track 31.21).

LiSnare 1

Inserts: UAD *Cambridge EQ*.

LiSnare 1 (Track 31.22) plays during the first part of the lead-in section and during the last four bars before the outro. This track was treated using the *Cambridge EQ*, where an HPF [168 Hz, 6 dB/oct] rolls off muddling lows (Track 31.23). Also, a high-shelving filter [6.32 kHz, 4 dB] brightens the snare and tunes it to the frequency spectrum (Track 31.24).

SnareB

Inserts: PSP *MasterQ*.

SnareB (Track 31.25) accompanies every other hit of LinSnare 1 during the first part of the lead-in section. It was treated with the *MasterQ*, where an HPF [74 Hz, 12 dB/oct] rolls off dispensable lows (Track 31.26) and a high-shelving filter [10 kHz, 4 dB] adds some extra definition (Track 31.27).

Claps

Inserts: t.c. electronic *ClassicVerb*.

There are five clap hits during the first part of the lead-in section, two of which conclude this first part. Only a reverb was added to these claps and the plugin was added as an insert. The wet/dry control was set to 50/50 and the pre-delay to 91 ms, and the early reflections were muted (Tracks 31.28 and 31.29).

LiMidSnare

Inserts: Logic *Channel EQ*.

Two hits of LiMidSnare play right before the second part of the lead-in section (along with two claps). The raw track (Track 31.30) was treated with the *Channel EQ*, where an HPF [130 Hz, 24 dB/oct] rolls off dispensable lows (Track 31.31) and a parametric filter [7.7 kHz, 5.5 dB, *Q* 1.9] adds brightness and definition (Track 31.32).

LiSnare 2

Inserts: UAD *Cambridge EQ*.

LiSnare 2 (Track 31.33) plays during the second part of the lead-in section. Its tonality was altered quite noticeably by the *Cambridge EQ*, mainly due to the HPF [428 Hz, 12 dB/oct], which dries out all the lows (Track 31.34). As part of frequency tuning, a parametric filter [900 Hz, 4.2 dB, *Q* 2] was also employed (Track 31.35) along with a high-shelving filter [6.32 kHz, 1.6 dB] (Track 31.36). At some point during the mix, a compressor was also loaded onto this track, but a punchy snare didn't quite work in the context of the relaxed lead-in section. What's more, the beat is not the most important element of this section—the real action occurs later.

Kick 2

Inserts: Sonnox *Oxford EQ*.

Kick 2 plays during the main and the straight sections of the production and it provides a layer to the main kick (Tracks 31.37 and 31.38). The only plugin inserted on this track is the *Oxford EQ*, where a parametric filter [95 Hz, –2.88, *Q* 1.4] attenuates disturbing content around 95 Hz (Tracks 31.39 and 31.40).

Kick WoodKnock

Kick WoodKnock only plays during the first main and the straight sections. It is a complementary layer for the kick and it was untreated (Tracks 31.41–31.43).

Main Snare

Inserts: UAD *Cambridge EQ*.

Main Snare appears during the main sections. In many productions such as "Donna Pomini," the kick is the mighty beat element and the snare gets much less stage space. To some extent, the Main Snare was mixed with slightly more entity than what some consider normal. The raw track (Track 31.44) was only treated with the *Cambridge EQ*, where a high-shelving filter [4 kHz, 1 dB] adds a touch of presence (Track 31.45). The level of this track was automated to rise during the second main section, but falls again in its second part.

Rvs Snare

Inserts: PSP *MasterComp*.

Rvs Snare plays during the first part of the first main section and throughout most of the second main section. It accompanies the main snare as a supportive beat element (Tracks 31.46–31.48). Rvs Snare is only treated with the *MasterComp* [threshold –16.1 dB, ratio 1.4:1, RMS sensing, hard-knee, attack 3.16 ms, release 1 second], which condenses the snare's dynamics (Track 31.49).

Straight Snare

Inserts: Sonnox *Oxford Dynamics*, Sonnox *Oxford EQ*, Logic *BitCrusher*.

Sends: SnareVerb (Sonnox *Oxford Reverb*).

Straight Snare (Track 31.50), which plays during the straight section, was first treated with the *Oxford Dynamics*, where both the compressor [threshold –22.5 dB, ratio 2:1, hard-knee, attack 15.3 ms, release 5 ms] and the gate [threshold –30 dB, range –40 dB, attack 5 µs, hold 550 ms, release 19.6 ms] are employed to add weight and punch (Tracks 31.51 and 31.52). Then, the *Oxford EQ* shapes quite drastically the sound of the snare using four bands: an HPF [271 Hz, 24 dB/oct] rolls off the body and dispensable lows (Track 31.53), a parametric filter [487 Hz, –8.44 dB, *Q* 2.83] attenuates disturbing low-mids (Track 31.54), another parametric filter [3.13 kHz, 7.71 dB, *Q* 2] adds snap (Track 31.55), and a high-shelving filter [3.32 kHz, 2.18 dB] adds some brightness (Track 31.56). The EQ is followed by the *BitCrusher*, which adds its own touch of grit [8 bits, drive 4.5 dB] (Track 31.57).

The snare was sent to a dedicated reverb—a modified version of the "Hall Full" preset on the *Oxford Reverb*. The send level to the reverb was automated so as to open on every other hit (with one hit being an exception). Apart from making the effect appear on every other hit, this also makes one hit appear at the back, and the other in front (Track 31.58). There is no practical purpose for this automation and it is hardly creative, but we do these things sometimes.

Whip Snare

Inserts: Sonnox *Oxford Dynamics*, Sonnox *Oxford EQ*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Whip Snare (Track 31.59) plays during the last four bars before the outro. It first goes through the *Oxford Dynamics*, where much like the Straight Snare both the compressor [threshold –17.7 dB, ratio 5:1, soft-knee, attack 16.8 ms, release 130 ms] and the gate [threshold … 17.2 dB, range –40 dB, attack 5 µs, hold 120 ms, release 20.9 ms] add snap and punch (Tracks 31.60 and 31.61). On the subsequent *Oxford EQ*, an HPF [210 Hz, 12 dB/oct] rolls off dispensable lows (Track 31.62) and a parametric filter [3.29 kHz, 3.2 dB, *Q* 1.96] emphasizes some edge (Track 31.63). Whip Snare is sent at high level to the ambiance reverb, which places it quite deep in the depth field and adds an apparent tail (Track 31.64).

March Snare

Inserts: SoundHack *+Chebyshev*, Logic *Gain*, UAD *GateComp*, t.c. electronic *Filteroid*.

March Snare is part of the buildup during the ocean section. The raw track already contained some room sound (Track 31.65). It is first treated with the *+Chebyshev*, which adds some harmonics that fatten the sound (Track 31.66). Then the *Gain* plugin was inserted for level automation, where the snare progressively rises in level (Track 31.67). The succeeding *GateComp* adds weight to the snare with its compressor facility [threshold –28.8 dB, ratio 4.64:1, slow attack, fast release] (Track 31.68). The *Gain* plugin was intentionally inserted before the *GateComp*, so the compression effect increases as the snare rises in level. On the succeeding *Filteroid*, a non-resonant HPF [185 Hz] and an LPF [17.8 kHz] contain the frequency extremes of the snare (Track 31.69).

Main hats

Inserts: Logic *BitCrusher*, SoundHack *+Chebyshev*, t.c. electronic *EQSat*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

The main hats are the only hats playing during the lead-in section and, as importantly, during the second part of the first break. They also play along other hats during the main sections. The clean sound of the raw track (Track 31.70) could work in the context of many similar productions, but I personally felt that a slightly more dirty sound would better suit the aggressive nature of this track. To achieve this dirty sound, two plugins were used. First, the *BitCrusher* reduces the bit resolution to 4 bits, which adds a noticeable amount of noise. In addition, 13 dB of drive causes clip distortion that adds healthy grit (Track 31.71). Then, the *+Chebyshev* adds some harmonics that enhance the effect (Track 31.72). The *BitCrusher* and *+Chebyshev* made the hats appear brittle and therefore detached from the mix. To ease this excess of highs, a high-shelving filter [5 kHz, –1.5 dB] was engaged on the *EQSat* (Track 31.73). The level of the main hats was automated with respect to the different sections in which they play. It might be worth noting that during the lead-in section, the hats are fairly visible as they play an important rhythmic role. They are also sent to the ambiance reverb, which moves them back from the front of the mix (Track 31.74).

Hats 2

Inserts: Logic *BitCrusher*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Hats 2 (Track 31.75) are the only hats playing during the straight section, but they play along other hats tracks during the first part of the second main section. As with the main hats, Hats 2 were treated with the *BitCrusher* [3 bits, 6 dB drive], which distorts them quite noticeably (Track 31.76). They were also sent to the ambiance reverb, which sends them back in the depth field (Track 31.77).

Hats 4 hits

Four hi-hat hits mark the border between the two parts of the first main section. These hits were left untreated (Track 31.78).

Hats Backbeat

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Hats Backbeat (Track 31.79) plays along with the main hats during the main sections. This track was untreated, but was sent to the ambiance reverb for depth positioning (Track 31.80).

Crash 1, Crash 2, and Crash 3

"Donna Pomini" involves three crash tracks. None of them were treated or sent to the reverb, but the levels of Crash 2 and 3 were automated between various hits.

WoodZest

Inserts: Logic *Channel EQ*.

WoodZest is part of the buildup to the first main section and plays for most of the production from that point onward (Tracks 31.81–31.83). It was only treated with an HPF [240 Hz, 24 dB/oct] from the *Channel EQ* (Track 31.84). Its level was automated; for example, to progressively rise during the ocean section.

PaperLoop

Inserts: Smartelectronix *Cyanide 2*, SoundHack *+Chebyshev*, Sonnox *Oxford EQ*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*), *Loop Delay* (PSP *84*).

PaperLoop plays during the straight and the second main sections. The raw track (Track 31.85) was distorted using *Cyanide* (Track 31.86). Then additional distortion was produced by the *+Chebyshev* (Track 31.87). The *Oxford EQ* was employed to tune this track into the frequency spectrum—a low-shelving filter [114 Hz,–6.27 dB] attenuates some lows (Track 31.88).

PaperLoop was sent at low level to the ambiance reverb in order to move slightly backward in the mix (Track 31.89). It was also sent to PSP *84*, with the send level automated to open during the last few hits before the ocean section. The delay time was set to quarter-note, –4.89 dB feedback gain, a resonant LPF in the feedback loop, and gentle modulation of the delay time (Track 31.90).

Bat Zips, Bat Zaps, Bat Hit1, Bat Hit2, Bat Hit3, Bat Iron, and Bat Wiper

Inserts: Smartelectronix *Cyanide 2*.

Sends (for some): Ambiance Reverb (t.c. electronic *ClassicVerb*).

Bat Zips, Bat Zaps, Bat Hit1, Bat Hit2, Bat Hit3, Bat Iron, and Bat Wiper (Tracks 31.91–31.98) were all processed using the same *Cyanide* configuration (Track 31.99). This produced a very distinctive sound, which adds some aggressiveness. Bat Zips, Bat Zaps, and Bat Iron were sent to the ambiance reverb in order to shift them backward in the mix (Track 31.100).

Beat Group

Inserts: PSP *Vintage Warmer*.

Although entitled "Beat Group," only four beat elements were routed to this audio group: Kick 2, Kick WoodKnock, Rvs Snare, and Hats Backbeat. On the group track, the *Vintage Warmer* was loaded and configured for parallel compression [mix 32.6 percent dry, drive 16.5 dB, knee 0, speed 51.5 percent, release 0.25]. The idea was to create a thicker sound layer and to draw some compression effect (Tracks 31.101 and 31.102).

Sound FX

Splutter

Inserts: Logic *BitCrusher*, Logic *PlatinumVerb*.

Splutter appears once, right before the introduction of the lead during the first break (Track 31.103). It was first processed using the *BitCrusher*, but the bit reduction facility was not utilized and the plugin was used purely for distortion and downsampling purposes [24 bits, drive 22.5 dB, downsampling 2×, clip level –7.8 dB] (Track 31.104). The *BitCrusher* is followed by the *PlatinumVerb*, which was inserted [73 percent dry, 49 percent wet] to send Splutter backward and add some effect to the dry sound (Track 31.105).

Laser Gun

Inserts: Logic *Channel EQ*.

Laser Gun appears right before the drop to the main section (Track 31.106). It was only treated by the *Channel EQ*, where an HPF [240 Hz, 24 dB/oct] rolls off dispensable lows (Track 31.107).

ChemiWind

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

ChemiWind gels the transition between the straight and the ocean sections (Track 31.108). It was only sent at high level to the ambiance reverb, which adds some dimension to it and generates a transition tail (Track 31.109).

FallBall

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

FallBall creates some tension toward the end of the first main section and the pitch dive leads into the straight section (Track 31.110). It also appears briefly between the two parts of the second main section. This track was only sent at high level to the ambiance reverb, mostly in order to place it behind the beat and the leads (Track 31.111).

SciFi Waves

SciFi Waves (Track 31.112) appear right before the first break, and also make a short appearance early in the first main section. This track was not treated, but its level was automated to rise before its appearance in the main section.

Bass

JoviBass

Two bass tracks were included on the multitrack. JoviBass (Track 31.113) plays up until the first break. This track was not treated.

ZugBass

In contrast to JoviBass (no treatment), ZugBass (Track 31.114), which plays from the first main section onward, was hopeless. Neither its frequency content nor its dynamics were workable and, despite some time and effort, it sounded wimpy, undefined, and flat. So ZugBass was excluded from the mix and I copied its source MIDI track into Logic and synthesized a new bass using the Access *Virus*. This new bass was called ZugBassR.

ZugBassR

Inserts: Sonnox *Oxford EQ*.

ZugBassR (Track 31.115) was equalized using the *Oxford EQ*, where an HPF [44.3 Hz, 36 dB/oct] sets the low-frequency boundary of the bass; this was done to prevent an excess of subs that occur on the low notes (Track 31.116). In addition, there is also a parametric filter [267 Hz, –4.8, *Q* 2.83] that combats an excess of low-mids (Track 31.117).

Vocal

Vox

Inserts: Logic *Vocal Transformer*, Logic *Channel EQ*, Logic *Guitar Amp Pro*.

The vocal on "Donna Pomini" (Track 31.118) called for a distinctive effect—just having a normal speaking voice would be mundane and characterless. Also, what's said, and the way in which it is said, suggests an abstract and somewhat dreamy message. A nice vocal effect would improve this mood. The first plugin in the inserts chain is the *Vocal Transformer*—it does not fulfill its original purpose, but adds a unique effect, which is a blend between a robotic hollowness and a wired distortion (Track 31.119). Next is the *Channel EQ*, where an HPF [450 Hz, 24 dB/oct] rolls off an excess of lows that really contributed nothing to the speaking voice (Track 31.120). The most dominant contributor to the final vocal effect is the *Guitar Amp Pro*, where, in addition to the guitar amp simulation, both the tremolo and the spring reverb are enabled (Track 31.121). The vocal track was sent to the ambiance reverb so as to shift it backward in the depth field and to give it some stereo size (Track 31.122).

The squeaking loop that appears in the second part of the main break and in the main sections has its origins in a saturated feedback loop of an automated delay on the words "each other." This loop was already bounced on the vocal track on the multitrack.

Other elements

Gidi FM

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Gidi FM plays during the lead-in section (Track 31.123) and also provides the FM riser during the breaks buildup (Track 31.124). This track was automated in level between and within the various sections. It is sent to the ambiance reverb, which pans it backward and increases its stereo size (Track 31.125).

Wacko

Inserts: Sonnox *Oxford EQ*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Wacko opens the production and also plays during the first main section and the straight (Track 31.126). Only a stereo version with imprinted delays was included on the multitrack.

This track was automated in three principal ways: its level was automated, for example, to drop twice during the intro section; while being panned center up until the straight section, it then shifts to the right [+21] to reduce its masking interaction with Straight Bass; for the same reason, the *Oxford EQ* is enabled at the beginning of the straight section and an HPF [149 Hz, 18 dB/oct] rolls off some colliding lows (Track 31.127). It was also sent to the ambiance reverb in order to blend it into the mix space (Track 31.128).

Wacko Filtered

Wacko Filtered plays for the first four bars of the production. On the multitrack, this track already contained the filter and the level automation (Track 31.129). It was not treated.

Wacko Support

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Wacko Support plays for two bars during the intro (Track 31.130). This track was mixed at a low level underneath Wacko and its level was automated to rise progressively. It was sent at a high level to the ambiance reverb, which sends it deep to the back of the sound stage (Track 31.131).

Drops

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Drops is one of the main elements playing during the lead-in section. The multitrack included both a version with no delay (Track 31.132) and a stereo version with a delay (Track 31.133). From the early mix stages to the final ones, the stereo version presented no issues, so it is the one used in the final mix. Drops is sent to the ambiance reverb, just to gel it to the mix ambiance (Track 31.134).

Drops Noise

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Just like with Drops, between the versions of Drops Noise with and without delay the former was used (Tracks 31.135 and 31.136). It is also sent to the ambiance reverb (Track 31.137).

SinSeqRoot, SinSeq3rd, SinSeq5th

Inserts: PSP *Nitro*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

SinSeqRoot, SinSeq3rd, and SinSeq5th all constitute a sequence of sine beats (Tracks 31.138–31.141). Each of the individual tracks was treated with the *Nitro*, which was configured for an autopan effect. The modulation frequency on each of the three tracks is different, which creates an unpredictable panning relationship between each track (Track 31.142). To detach the sequence from the front of the mix, each track was also sent to the ambiance reverb (Track 31.143).

SinSeq group

Inserts: PSP *608 MultiDelay*.

The three SinSeq tracks are all routed to an audio group. The inserted *608* is set for 35 percent dry and 65 percent wet and two taps are operational—one with a delay of 372 ms and the other with 743 ms. The feedback control on the second tap is set to 32.5 percent. This addition of delays softens the SinSeq and makes it somewhat more dreamy and distant (Track 31.144). The level of the SinSeq group was automated.

LeadSeq

Inserts: UAD *Fairchild*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

LeadSeq (Track 31.145) is introduced during the second half of the first break and is one of the most important elements in the production. This lead occupies a huge amount of space on the frequency spectrum and as it clearly involves unison it also takes up quite some space on the stereo panorama. It is somewhat surprising that this lead fits nicely into the mix without any frequency or stereo treatment. The *Fairchild* inserted on this track was employed to condense the level of this lead and add some more weight—essentially making the lead even bigger (Track 31.146). The level of this track was automated; for example, it is made louder on the first main section. It is also sent at low level to the ambiance reverb, which gels it into the mix ambiance and shifts it slightly backward (Track 31.147).

LoFi Saw

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

LoFi Saw accompanies LeadSeq during the first main section and the first part of the second main section. Between the versions on the multitrack with and without delay, I chose to use the one with delay (Tracks 31.148 and 31.149). This track was untreated,

but it was sent at a high level to the ambiance reverb. This not only places it behind LeadSeq, but also fills the ambiance during the most powerful moments of the production, making it denser and richer (Track 31.150).

Straight Bass

Inserts: UAD *Cambridge EQ.*

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Straight Bass (Track 31.151) plays during the straight and ocean sections. The *Cambridge* inserted on this track only has one parametric filter engaged [90.8 Hz, –2.4 dB, *Q* 1.66], which eases an excess of lows around 90 Hz that muddled the mix (Track 31.152). This track was also sent to the ambiance reverb in order to send it slightly backward and gel it to the mix ambiance (Track 31.153).

Straight Beep

Straight Beep plays during the straight section only. I wanted the beep to appear distant and to audibly excite this reverb or another. Initially, I used the dry version from the multitrack and sent it at a high level to the ambiance reverb (Tracks 31.154 and 31.155). However, the multitrack also contained a version with a reverb and auditioning it at a later mix stage worked out better (Track 31.156). It is the only track from the multitrack that had a version with a reverb that was used in the final mix.

Leadar

Inserts: PSP *MasterComp.*

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Leadar plays from the beginning of the ocean section until the end of the production. Between the versions on the multitrack with and without delay, I chose to use the former (Tracks 31.157 and 31.158). This track was compressed using the *MasterComp* [threshold –12.93 dB, ratio 1:30:1, RMS sensing, hard-knee, attack 10 µs, release 100 ms], which thickens it and makes it bigger (Track 31.159). It was also sent to the ambiance reverb, which gels it to the sound stage but also makes its stereo appearance richer and bigger (Track 31.160).

OceanPad

Inserts: PSP *MasterQ.*

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

OceanPad plays during the ocean section (Track 31.161). This track was compressed with the *MasterQ*, where a high-shelving filter [6.41 kHz, –4.15 dB, *Q* 0.5] rectifies an excess of brittle highs (Track 31.162). The ocean section is also some kind of a break that I felt needed to have an extra dreamy and spacious sound. OceanPad was sent at an extremely high level to the ambiance reverb, which makes it appear distant and facilitates the required dreamy sensation (Track 31.163).

OceanSup

Inserts: PSP *84*.

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

OceanSup is a complementary track to OceanPad—while OceanPad is more concerned with the soundscape, OceanSup gives a clearer idea of the melody being played (Tracks 31.164 and 31.165). The raw track (Track 31.166) was first treated with the PSP *84* [30 ms delay on the left channel only, no feedback], which combats this track's presence in the busy center—the delay converts its central image to a sound bound to the stereo extremes (Track 31.167). It was also sent at a high level to the ambiance reverb to help project the wistful ambience that this section demanded (Track 31.168).

LowHarm, HighHarm

Sends: Ambiance Reverb (t.c. electronic *ClassicVerb*).

Both LowHarm and HighHarm only appear at the end of the ocean section (Tracks 31.169 and 31.170). Both tracks were sent at a high level to the ambiance reverb in order to place them far away in the depth field (Track 31.171). Both tracks were mixed at a low level, which means that rather than being clearly present they blend with the overall atmosphere.

ContraWaves

ContraWaves appears during the eight outro bars as a support layer to Leadar (Track 31.172). This track was not treated.

32 "The Hustle" (DnB)

"The Hustle" is a powerful DnB production. The 41 tracks submitted with this project—mostly the outcome of original programming—constitute a meticulous arrangement that could easily serve as a production masterclass. Indeed, from a mix engineer's point of view, it was exciting to have an insight into how this epic tune was constructed.

Drum 'n' bass productions, as the name suggests, are characterized by dominant beats and basslines. In mixes, the two tend to be loud to an extent that would not be considered sensibly balanced in many other genres. In addition, most mixes also entail a deep and wide sound stage and a pronounced sense of space—all the outcome of evident reverbs. The fast tempo of DnB productions (174 bpm in the case of "The Hustle") and the typically dense arrangements mean that mostly we are trying to find space in the mix for the various elements.

While many DnB productions are heavily based on a massive beat and a grandiose bassline, the qualities of "The Hustle" are distributed between other elements as well. I thought that the mix should reflect this, so the general approach in mixing the various elements was slightly more balanced compared with many traditional DnB mixes.

It is not uncommon for a rough mix to turn into the mix itself; such was the case with this production. By the time I had studied the various tracks, sections, and nature of the production, I already had a mix-in-progress, and there was little point starting all over again.

Track 32.1 is the producer's rough mix; Track 32.2 is the mix-ready version with rough levels and panning but no processors or effects; Track 32.3 is a snapshot of the mix-in-progress; Track 32.4 is the final unmastered mix; and Track 32.5 is the mastered version.

Ambiance reverb

Audio Ease *Altiverb*.

The reverb choice in productions such as "The Hustle" is a crucial aspect of the mix. In such dense arrangements, any instrument that doesn't need to be forward is panned backward—a panning mostly achieved using reverbs. In addition, the typical prominent space of DnB mixes means that the reverbs used dictate much of the mix character. The plan for this mix was to use one main reverb to which many instruments are sent. This reverb is used for frontback panning and contributes to the space attributes of the mix.

A "draft" reverb was used during the early mix stages (a common practice), and the choice of the final ambiance reverb was made once enough tracks had been sent to that reverb and once the mix had progressed enough to hint at the direction the ambiance needed. After a little experimentation, I chose the Audio Ease *Altiverb*, its "Amsterdam Concertgebouw" IR collection (under the "Concert Halls Large" category), and the "Stereo to Stereo far Omnis" IR. This specific IR was chosen as it excelled at depth positioning, as it handled transients effortlessly, and as its warm, defined-yet-elegant nature blended superbly into the mix. The tail of this ambiance reverb can be heard on Tracks 32.6 and 32.7, while Tracks 32.8 and 32.9 provide a comparison between a mix with and without the ambiance reverb. One issue with the ambiance reverb was that it felt too wide, so the *Altiverb* output was narrowed by the "MOTU Trim" (Tracks 32.10 and 32.11).

Drums

Attack Kick

Inserts: Sonnox *Oxford EQ*.

Sends: Kick SC.

Attack Kick (Track 32.12) is the main kick track. It was treated with the *Oxford EQ*, where a low-shelving filter [136 Hz, 8.2 dB] adds a healthy lows impact (Track 32.13). A steep HPF [42 Hz, 30 dB/oct] prevents boominess and filters the excess of very low frequencies, which many domestic systems cannot reproduce anyway (Track 32.14). There is also a parametric filter [273 Hz, –9 dB, *Q* 1.5] removing a low-mids tap (Track 32.15). Attack Kick is sent to a bus, which sources the side-chain of the sub-bass gate.

Dirt Kick

Inserts: MOTU *Trim*.

The Attack Kick is layered with Dirt Kick (Track 32.16), which adds some oomph and noise that contributes some definition (Tracks 32.17 and 32.18). Despite the fact that both are samples, the Attack Kick and the Dirt Kick were out of phase with each other, so the phase of the Dirt Kick was inverted using the MOTU *Trim* plugin (Tracks 32.19–32.22); otherwise it is untreated.

Kick triggers

Virtual Instrument: MOTU *Model 12.*

Still unhappy with the sound of the kick, I decided to add another layer of samples. This was achieved by exporting the kick's MIDI track from the original Logic session, importing it into Digital Performer, and routing it to the MOTU *Model 12* Drum Module. The kick061 was chosen from more than 100 kick samples that ship with *Model 12* (Track 32.23). In addition, I programmed a MIDI sequence with a kick on each down-beat and sent it to the same *Model 12*, this time triggering the kick091 sample. This additional trigger was mixed so it only adds a gentle accent to the downbeats (Track 32.25). Finally, I added sub-bass to the kick by gating a 50 Hz sine wave with respect to the original kick track (Track 32.26).

Thump Kick

There are only six Thump Kick (Track 32.27) hits in the arrangement, three in each break. This track was untreated.

Snare

Inserts: MOTU *MasterWorks EQ*, Sonnox *Oxford Dynamics.*

Sends: Snare Reverb (UAD *Plate 140*).

One problem with the main snare track (Track 32.28) was that it was too bright and did not sit well in the frequency spectrum—it seemed to end up on higher frequencies than it should have done. To correct this, a high-shelving filter [4 kHz, –1.6 dB] on the *Master-Works EQ* was employed to soften the highs (Track 32.29). Following this, the snare sat better in the frequency spectrum but it lost some presence. In order to return it to the limelight, the compressor [threshold –10.2 dB, 2:1 ratio, hard-knee, attack 28.6 ms, release 130 ms] of the *Oxford Dynamics* was employed to add some snap (Track 32.30).

I also experimented with adding a reverb to the snare using the *Plate 140* (Track 32.32.). It is the nature of reverbs to soften sounds, so in order to achieve some thrust I gated the reverb using the MOTU dynamics (Track 32.32). Since the reverb smears the snare image and sends it slightly backward, it is only mixed during the intro, where the arrangement is sparser and more relaxed.

Sub Snare

Sub Snare (Track 32.33), which was left untreated, is complementary to the snare track that was layered quietly underneath it (Tracks 32.34 and 32.35).

Pend Snare

Inserts: MOTU *Gate*, MOTU *Parametric EQ*, Sonnox *Oxford Dynamics.*

Pend Snare (Track 32.36) was one track that I felt contributed little to the overall snare sound, so instead of layering it with the snare track, it was mute-automated to only appear when the snare plays eighth-notes, therefore accenting the momentary

intensity of the beat (Tracks 32.40 and 32.41). The raw track involved a very long decay, so to make it punchier it was gated using the MOTU *Gate* (Track 32.37). The gate is followed by the *Parametric EQ*, where a high-shelving filter [1.3 kHz, 4.7 dB] brightens what was a dull sound (Track 32.38). To add even more punch, the *Oxford Dynamics* compressor [threshold … 18 dB, ratio 3.17:1, attack 10.6 ms, release 5 ms, soft-knee 20 dB span] accents the snare's attack (Track 32.39).

Hi Hats

Inserts: MOTU *AutoPan*, MOTU *MasterWorks EQ*, MOTU *Dynamics*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

In addition to the kick and snare, the main beat also consists of the tracks Hi Hats, Comp Break, and Tamb. The Hi Hats (Track 32.42) go through the *AutoPan*, which adds movement. As the hi-hats shift between left and right, their perceived level rises and falls (partly due to masking), which adds another element of cyclic variation. In order to make the effect not too obvious, the depth on the *AutoPan* was set to 43 percent; the modulation is synced to a bar (Track 32.43). The *AutoPan* is followed by the *MasterWorks EQ*, where a high-shelving filter [3.5 kHz, 2.8 dB] adds definition and spark (Track 32.44). Although not everyone would agree, I felt that brightening the hi-hats further (Track 32.45) would detach them from the mix and cause exaggerated frequency separation. Following the *MasterWorks EQ* is the MOTU *Dynamics*, whose role is explained in the next section. One of the problems with the hi-hats was that they sounded too in front and as such a bit like a fly buzzing in the front of the mix. Sending them to the ambiance reverb shifted them backward and rectified this buzzing sound (Track 32.46).

Comp Break

Inserts: MOTU *Parametric EQ*, MOTU *Dynamics*.

Comp Break (Track 32.47) is another important component of the beat. Keeping the kick and snare in this loop served little purpose since each drum has respective tracks already. Using the *Parametric EQ*, an HPF [725 Hz] filters much of the kick and snare, while also clearing some low-mids space for other tracks (Track 32.48). As part of tuning this track to the frequency spectrum and in order to give it more presence, a high-shelving filter [5 kHz, 5.7 dB] was also applied (Track 32.49). The Comp Break was panned off-center to the left [<41] so as to ease masking interaction with the kick and snare.

Later on in the mix, I wanted the leads to stand out more. One way to emphasize them was by ducking other instruments in relation to them—as if they are so powerful that they bring down other instruments. Ducking the whole beat would also attenuate the important kick and snare, so instead only Hi Hats and Comp Break were ducked. This was achieved by inserting the MOTU *Dynamics* compressor on the two tracks with each sidechain set to the same bus. The bus was sourced from a send on the Dina lead track—a send that was pre-fader, so level alterations to the lead would not affect the ducking. The major setting on the compressors was the release, as it determined how quickly Hi Hats and Comp Break would recover after being ducked. The final settings were 560 ms for the Comp Break and 480 ms for the Hi Hats (Tracks 32.50 and 32.51).

Tamb

Inserts: MOTU *Parametric EQ.*

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

The Tamb (Track 32.52) is introduced in the arrangement after the first break. It is mixed, so it blends with the Hi Hats and Comp Break rather than being a distinctive sound (Tracks 32.56 and 32.57). An HPF [1.3 kHz] on the MOTU *Parametric EQ* rolls off dispensable frequency content, which again clears space on the low-mids (Track 32.53). In addition, a parametric filter [5.3 kHz, 3.3 dB] fills some missing frequency space (Track 32.54). The Tamb is also sent to the ambiance reverb, which blends it into the depth field (Track 32.55). The Tamb was panned slightly off-center to the right [>12], a position in which it best escaped masking without shifting too far (from the beat) to the extreme. It is not ducked since it is masked anyway by the leads, at the moments when it would have been ducked.

Break Beat

Inserts: MOTU *Delay*, SoundHack *+Chebyshev*.

The drums during the break are based on four tracks: Break Beat, Hip Hop Beat, Claps, and Half Time. Break Beat (Track 32.58) and Hip Hop Beat are two competing tracks and I wanted to maintain separation between them—had the two been blended around the center, the details of both would have been lost and the two would have combined unimpressively. A break, by nature, involves a departure from the main theme, and this applies to the way we mix it as well. So, while Hip Hop Beat was panned center, Break Beat goes through the MOTU *Delay*, which only delays the right channel by a quarter-note (Track 32.59). This creates an engaging effect that separates the two competing tracks. Also, this means that during the break some intensity shifts to the extremes and gives an overall wider stereo impression compared with the sections before and after it. The Break Beat delay is followed by the *+Chebyshev*, which adds a touch of distortion (Track 32.60).

Hip Hop Beat

Inserts: Sonnox *Oxford Transient Modulator*.

The raw Hip Hop Beat track (Track 32.61) could benefit from more vibrant dynamics, so the *Transient Modulator* [ratio 0.12] was employed to accent the transients on this loop (Track 32.62).

Hip Hop Beat distorted layer

Inserts: UAD *Preflex*.

The raw Hip Hop Beat track was duplicated and the new track was distorted using the *Preflex*. The distorted track is layered with the original Hip Hop Beat and automated to progressively rise in level (Tracks 32.63 and 32.64).

Claps

Inserts: PSP *MasterQ*.

Claps (Track 32.65) were only treated with the *MasterQ*, where an HPF [1.8 kHz] rolls off frequency content that contributes little to their sound (Track 32.66).

Half Time

Inserts: MOTU *MasterWorks EQ*, Sonnox *Oxford Dynamics*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*), *Closing Delay* (MOTU *Delay*).

Half Time (Track 32.67) is a percussive loop that comes and goes throughout the production. The untreated track contained excessive lows and low-mids noise that contributed nothing but reduced clarity. So it was first filtered by an HPF [214 Hz, 24 dB/oct] from the *MasterWorks EQ* (Track 32..68). Also, a parametric filter [1.2 kHz, 6.4 dB, *Q* 0.85] added some flesh to what I found to be thin (Track 32.69). Following the equalizer, an *Oxford Dynamics* was inserted, where the compressor [threshold –30 dB, ratio 2:1, attack 2.2 ms, release 52 ms] condensed the levels of the per-cussive sounds, adding some powerful impact (Track 32.70). Half Time was also sent to the ambiance reverb in order to place it backward in the depth field (Track 32.71). The level of this track was automated, where it became louder [2.7 dB] during the second break. In addition, a send to a ping-pong delay was automated to take effect during the last bar of the outro. The MOTU *Delay* was set to quarter-note delays with a feedback LPF at 3 kHz (Track 32.72). This delay appears to escape to the left speaker earlier than normal with this effect, and, although this could have been corrected, I decided to keep it.

Bongos 1 and Bongos 2

Inserts: MOTU *MasterWorks EQ*.

Sends: Bongos Reverb (UAD *DreamVerb*).

Bongos 1 and Bongos 2 (Track 32.73) were both treated with a high-shelving filter [4 kHz, 6.8 dB] using the *MasterWorks EQ* (Track 32.74). Mostly, this was done in order to increase their definition. It was hard not to imagine the bongos with a nice, clear reverb. But the ambiance reverb was not fit for the job—the bongos sounded either too dry or too far away (Track 32.75). Also, in the specific case of the bongos, I was after brilliance rather than the native warmth of the ambiance reverb. So the bongos were given their own reverb—the UAD *DreamVerb* and its "Church Choir" preset (Track 32.76). This preset was only modified by a tiny pull of the high-mids on the pre-reverb EQ.

I did not want the two bongo tracks to be too far apart, yet I did not want them panned to the busy center. So each track is panned slightly off-center toward a different extreme, just enough so the two can be distinguished on the horizontal plane. The level of Bongos 1 is automated to balance some hits that jumped out of the mix.

Pen Drums

Inserts: UAD DM-*1 L Delay Modulator*, t.c. electronic *Filteroid*.

On the rough mix submitted by the artist, Pen Drums (Track 32.77) were treated with heavy flanging. To re-create this effect, I used the *DM-1 L* (Track 32.78), which was followed by the *Filteroid* (Track 32.79). The *Filteroid* was set with a resonant LPF, with its cut-off frequency slightly modulated (synced to a quarter-note).

In the original arrangement, the Pen Drums appear a few times throughout the production. The problem was that only during the second break could they be clearly heard—in all other places, they were poorly defined (Track 32.80). This led me to mute all the instances of these drums except those in the second break (Track 32.81).

Crash

Inserts: MOTU *MasterWorks EQ*.

The Crash (Track 32.82) is only treated with a high-shelving pull [4 kHz, –2.8 dB] using the *MasterWorks EQ* (Track 32.83). This was done in order to tune it into the frequency spectrum, which was too bright.

Beat group

Main Group Inserts: PSP *Nitro* (automated).

Parallel Group Inserts: MOTU *Dynamic*.

Parallel compression was applied on the beat. All the beat tracks are routed to a bus, which feeds two auxiliary tracks. One auxiliary is the main beat group and is only processed for two bars during the second theme section with a flanger from the Nitro (Tracks 32.84 and 32.85). The same group is automated so the beat during the intro and breaks is 1.7 dB quieter than during the theme sections. The second auxiliary is the compressed layer, processed by the compressor [threshold –28 dB, ratio 1.3:1, attack 7.3 ms, release 10 ms] of the MOTO *Dynamics* (Tracks 32.86–32.89).

Motif elements

Five tracks constitute the bassline and leads: Sick as Funk, Bass, Focus Bass, Dina, and Worm.

Sick as Funk

Inserts: MOTU *MasterWorks EQ*, t.c. electronic *Chorus Delay*, MOTU *Trim*, MOTU *Quan Jr*.

Sick as Funk was the most problematic track in this production—it choked (Track 32.90). This created a sudden drop in level, which sounded more erroneous than intentional. The first processor Sick as Funk goes through is the *MasterWorks EQ*, where a low-shelving filter [107 Hz, –12 dB] reduces some excess of lows (Track 32.91)

and a high-shelving filter [4.3 Hz, 5.6 dB] brightens up what was a dull sound (Track 32.92).

The main plugin employed to conceal the choking sound is the t.c. electronic *Chorus Delay*, inserted after the EQ. The "Two-Track" preset was loaded on the *Chorus Delay*, and its 40 ms of delay and 20 percent feedback were sufficient to fill the quick level drop (Track 32.93). Another plugin that completed the concealment is the *Quan Jr*, which reduced the bit resolution to 4 bits (Track 32.94). This adds a noticeable distortion that also promotes the aggressive nature of the track. The *Quan Jr* only offers a reduction to whole bit values, as opposed to some bit-reduction plugins that offer further bit divisions such as 4.5 (while a half bit might not make theoretical sense, it is perfectly possible to have fractional bit reduction in the floating-point domain). When a bit-reduction process only offers whole values, it is possible to fine-tune the effect by boosting or attenuating the signal prior to the bit reduction. And so, the *Quan Jr* is preceded by the MOTU *Trim* plugin, which boosts the signal by 2.4 dB. The bit reduction introduced some high-frequency noise, which was eliminated by an LPF [13.2 kHz, 12 dB/oct] from another *MasterWorks EQ* that succeeded the *Quan Jr* (Track 32.95). On the same EQ, a parametric filter [4 kHz, 5.6 dB] adds some presence and definition (Track 32.96). Sick as Funk is panned slightly off-center to the left [<11].

Bass

The Bass track (Track 32.97) is layered with the Sick as Funk track to complement it. It was left untreated and panned hard-center (Tracks 32.98 and 32.99).

Focus Bass

Inserts: Sonnox *Oxford EQ*, MOTU *Dynamics*.

Focus Bass (Track 32.100) was treated with the *Oxford EQ*, where a parametric filter [132. Hz, –8.4 dB, *Q* 16] treated some resonance that occurred on the last note of a rising sequence (Track 32.101). Also, a low-shelving filter [89.4 Hz, –6.3 dB] reduced some overemphasized lows (Track 32.102). A high-shelving EQ pulls 4 dB at 9.9 kHz to reduce high-frequency noise (Track 32.103). The EQ was followed by the MOTU *Dynamics*, where a compressor is set to duck the bass with relation to the Dina track. The release was set to a rather long 3 seconds, so that around 5 dB of gain reduction recovers just before the next Dina hit (Tracks 32.104–32.107). Focus Bass was panned slightly off-center to the left [<12].

Dina

Inserts: Sonnox *Oxford Dynamic*.

Sends: Ducking Side-Chain.

Dina (Track 32.108) is the main lead track. It is only treated with the compressor [threshold –22.5 dB, ratio 4.26:1, hard-knee, attack 52 ms, release 5 ms] on the *Oxford Dynamics*, which was employed to add attack (Track 32.109). The Dina lead is the sole

track to duck other tracks such as the Focus Bass and Hi Hats. It is panned slightly off-center to the left [<16].

Worm

Worm (Track 32.110) complements the Dina track. It could have been panned to the same position as Dina (Track 32.111), but in order to open up the sound it was panned off-center to the right [>26] (Track 32.112). Otherwise, it is untreated.

Lead group

Main Group Inserts: Sonnox *Oxford EQ*.

Parallel Group Inserts: MOTU *PreAmp-1*, MOTU *Trim*, MOTU *MasterWorks EQ*.

In order to introduce some extra power and aggression, parallel distortion was applied on the five tracks constituting the bassline and leads. This was achieved by sending the five tracks to a bus, and feeding that bus into two aux tracks. One aux track is simply an audio group, and is only treated with the *Oxford EQ*, where a parametric filter [125 Hz, –2.88 dB, *Q* 1.87] tunes the leads to the frequency spectrum by reducing some excess content around the 125 Hz area (Track 32.113). The other aux contributes the distorted layer, where the *PreAmp-1* severely distorts the signal (Tracks 32.115 and 32.116). The output of the *PreAmp-1* involves distinct stereo spread that was unwanted. So it is followed by the MOTU *Trim*, which converts the stereo signal into mono (Track 32.117). The distortion added a generous amount of high-frequency noise, so the *Trim* plugin is followed by *MasterWorks EQ*, where a 12 dB/oct LPF rolls off at 14.1 kHz, and a high-shelving filter attenuates 3.6 dB at 4 kHz (Track 32.118).

Initially, the distorted layer was mixed 8.8 dB below the untreated audio group, but this resulted in saturated lows that lacked tightness (Track 32.119). So the balance between the two was adjusted, ending up with the distorted layer 14.3 dB below the audio group (Track 32.120). This means that the distorted layer ended up being felt more than clearly heard, but it still contributes some extra power (Track 32.121).

The level of the lead group was automated. First, the level rises just before the second break (Tracks 32.122 and 32.123). Then, during the second theme section, the level is ridden to accent single note hits (Tracks 32.124 and 32.125).

Hemorrhage

Inserts: MOTU *Delay*, MOTU *MS Decoder*, MOTU *Dynamics*.

Hemorrhage (Track 32.126) plays between the main motif lines and during the second break. The original mono track would have competed with the beat and other main elements had it been panned center, and would have created stereo imbalance had it been panned to either extremes. So, to make it stereophonic, the Haas trick was applied using 20 ms of delay on the right channel only using the MOTU *Delay* (Track 32.127). But the new stereophonic track was too wide, so the MOTU *MS Decoder* narrows its stereo width; this worked well despite the comb filtering that could have occurred between the delayed and nondelayed channels (Track 32.128). Then, the

MOTU *Dynamics* is used to duck Hemorrhage with relation to Dina (Tracks 32.129 and 32.130). The level of this track was also automated to dive by 3.6 dB during the second break. Hemorrhage is panned off-center to the right [>24].

Jumpy

Inserts: MOTU *MasterWorks EQ*, MOTU *Delay*, MOTU *Dynamics*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Jumpy is a complementary layer for Hemorrhage—it was mixed only to blend with Hemorrhage, not to compete with it; thus, it is not easily discerned (Tracks 32.131–32.132). It was first treated with the *MasterWorks EQ*, where an HPF [263 Hz, 24 dB/oct] rolls off dispensable lows. Also, a parametric filter [4.4 dB, 1.86 kHz, *Q* 0.64] adds some presence, and a high-shelving filter [8.4 dB, 9.16 kHz] pulls some high frequencies as part of frequency tuning (Tracks 32.133 and 32.134). Jumpy was a stereo track that did not blend well with the stereo panorama. So the MOTU *Delay* was used with eighth-note delay on the right channel only—no delay on the left channel. This widened the stereo image to the extremes and added some rhythmical effect (Track 32.135). Like Hemorrhage, Jumpy was also ducked in relation to Dina using the MOTU *Dynamics*. Jumpy was also sent to the ambiance reverb, which pans it slightly backward (Track 32.136).

Sick Warble

Inserts: MOTU *MasterWorks EQ*, MOTU *Delay*, MOTU *Dynamics*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Sick Warble is a variation of Jumpy. It only plays in the first theme section, where eight bars of Hemorrhage with Jumpy are followed by eight bars of Hemorrhage with Sick Warble. Like Jumpy, it is a complementary layer for Hemorrhage and was mixed to only blend underneath it (Tracks 32.137 and 32.138). On the *MasterWorks EQ*, an HPF [346.7 Hz, 18 db/oct] rolls off some dispensable lows (Tracks 32.139 and 32.140) and a parametric filter [7.6 dB, 3.24 kHz, *Q* 0.83] boosts some presence (Track 32.141). The MOTU *Delay* also adds eighth-note delay, only this time the non-delayed signal is panned hard-right and the delayed signal is panned hard-left (Track 32.142). The MOTU *Dynamics* ducks Sick Warble in relation to Dina. In addition, Sick Warble is also sent to the ambiance reverb for backward panning (Track 32.143).

Pads

Atoms 1

Inserts: Sonnox *Oxford EQ*, MOTU *Trim*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Atoms 1 (Track 32.144) plays during the introduction and during the second break. First, it goes through the *Oxford EQ*, where a parametric filter [1.3 kHz, –5.6 dB, *Q* 2.83] pulls

some edge to soften the sound (Track 32.145). In addition, a high-shelving filter [5.2 kHz, 4 dB] adds some brilliance and definition (Track 32.146). Atoms 1 contributes to the main melodic line during the introduction. Therefore, I found the autopan effect imprinted on the raw stereo track disturbing. To combat this, the stereo track was summed to mono using the MOTU *Trim* (Track 32.147). It was sent to the ambiance reverb to fabricate some atmosphere and to pan the pad backward from the beat. The ambiance reverb also compensates for the mono summing by adding some stereo width (Track 32.148).

Atoms 2

Atoms 2 (Track 32.149) was one track I found annoying, even after applying various treatments to it. I found its sharp character to contribute little to the overall mood of the production and decided to exclude it from the final mix.

Korg 174

Inserts: PSP *MasterQ*, MOTU *Delay*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

The Korg 174 (Track 32.150) only plays during the intro. Four bands on the *MasterQ* were employed to tune this track into the frequency spectrum: an HPF [145 Hz, 12 dB/ oct], a low-shelving filter [1 kHz, –8.2 dB], a parametric filter [2.27 kHz, –9.47 dB, *Q* 1.1], and a high-shelving filter [9.5 kHz, –2.87 dB] (Track 32.151). The original stereo track had a limited stereo width and I wanted this pad to stretch to the extremes. So the MOTU *Delay* was inserted to produce the Haas trick, with 30 ms delay on the right channel only (Track 32.152). This track is also sent to the ambiance reverb in order to detach it from the front of the mix and to blend it into the mix space (Track 32.153).

Horns and brass

Curtis Rev

Inserts: MOTU *Tremolo*, MOTU *Parametric EQ*, MOTU *Dynamics*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Curtis Rev (Track 32.154) is introduced in the intro, and then plays again during the theme sections. It was first treated with the *Tremolo* [bar-synced, 72 percent depth] in order to tighten it to the rhythm and to enhance its come-and-go nature (Track 32.155). The *Tremolo* was followed by the *Parametric EQ*, where an HPF [396 Hz] removes muddiness (Track 32.156) and a high-shelving filter [2.7 kHz, 4.5 dB] adds some brilliance and definition (Track 32.157). The MOTU *Dynamics* ducks the track in relation to Dina. This track is sent to the ambiance reverb for spatial positioning (Track 32.158).

Curtis Rev was automated in two principal ways. First, it was muted during the first theme section, which clears some space in the mix. Second, it was automated to rise in level during the second theme section, including a rise after the first eight bars of each Hemorrhage phrase.

Horns

Inserts: UAD *DM-1 Delay Modulator*, t.c. electronic *EQSat*, t.c. electronic *Filteroid*.

Horns (Track 32.159) is introduced during the intro and plays during the two breaks. The raw track included open-close hi-hats, which fortunately responded well to the applied equalization (by that I mean hardly limiting the processing). The original mono track is first treated with a chorus from the *DM-1*. The chorus turns this track into stereo, with the dry/wet control determining the stereo width (ending up at 27.5 percent dry). The chorus also adds some spatial dimension that sends the horns backward (Track 32.160). Next, a low-shelving filter [192 Hz, –5 dB] on the *EQSat* reduces some lows (Track 32.161) and a high-shelving filter [7 kHz, 3 dB] adds some brilliance and tunes this track into the mix spectrum (Track 32.162).

Horns is automated in two ways. First, its level is automated to rise during the intro, then it drops with the first break and starts rising again. Second, during the first part of the second break, a resonant LPF from the *Filteroid* sweeps from 668 Hz up to 20 kHz (Track 32.163).

Hip Hornz 2

Inserts: MOTU *MasterWorks EQ*, UAD *DM-1 Delay Modulator*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Hip Hornz 2 (Track 32.164) was mixed to appear distant, somewhat like distant ship horns. First, this track goes through the *MasterWorks EQ*, where an HPF [263 Hz, 18 dB/oct] removes an excess of lows that resulted in muddiness and lack of definition (Track 32.165); in addition, a high-shelving filter [1.3 kHz, 7.2 dB] adds some missing highs and extra definition (Track 32.166). Just like with Horns, the same chorus from the *DM-1* creates some stereo width and sends the track backward (Track 32.167). The track is sent at a high level to the ambiance reverb, which places it deep within the mix space and adds a nice, warm reverb tail (Track 32.168).

Hip Hornz 3

Inserts: MOTU *Parametric EQ*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Hip Hornz 3 (Track 32.169) is something of an answer to a question asked by Hip Hornz 2—I wanted them to appear closer, so as to create a contrast between the two. The track is only processed with the *Parametric EQ*, where an HPF [289 Hz] rolls off an excess of lows (Track 32.170). It is sent to the ambiance reverb but with less level than Hip Hornz 2; thus, it appears closer (Track 32.171).

Curtis Verb

Inserts: MOTU *MasterWorks EQ*, MOTU *Dynamics*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Curtis Verb (Track 32.172), essentially a burst of two trumpet-like notes with reverb, was first treated with *MasterWorks EQ*, where a high-shelving filter [4 kHz, –4 dB] tuned it to the frequency spectrum by softening its highs (Track 32.173). One issue with this track was that it appeared too much in front. It could have been sent backward with more reverb, but, as discussed below, the reverb added to this track plays a unique roll that restricted changes to the reverb send level. So the job of sending this track backward was assigned to the MOTU *Dynamics*, where a compressor [–15 dB threshold, 2:1 ratio, 0.1 ms attack, 100 ms release] contained its attack (Track 32.174).

Originally, an HPF on the *MasterWorks EQ* was employed to roll off the lows, but after sending this track to the ambiance reverb (Track 32.175), disabling the HPF resulted in a metallic thunder that I liked (Track 32.176). This effect, which was discovered by accident, is kept in the final mix.

Risers

Riser and Reversed Rise

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Riser (Track 32.177) and Reversed Rise (Track 32.178) mostly play together, but at certain points play alone. Neither track is treated, but Riser is sent to the ambiance reverb with a generous send level (Track 32.179).

Bent Phase

Inserts: MOTU *MasterWorks EQ*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Bent Phase (Track 32.180) was treated with the MOTU *MasterWorks EQ*, where an HPF [525 HZ, 18 dB/oct] removes dispensable low-mids and lows (Track 32.181). It is sent to the ambiance reverb for depth positioning (Track 32.182).

Hi Synth

Inserts: MOTU *PreAmp-1*, MOTU *Dynamics*, MOTU *Trim*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

One problematic track was the Hi Synth (Track 32.183) because it sounded synthetic and foreign to the arrangement. At certain points, this track was a candidate for exclusion, but it ended up lightly buried after some treatment. First, it was distorted using the *PreAmp-1*, which adds some grit (Track 32.184). Second, it was compressed using the MOTU *Dynamics* (Track 32.185)—the untreated track rises both in level and pitch, and I found that the rises in pitch were sufficient for the buildup, so the compressor contained the level rise. Finally, it goes through the MOTU *Trim*, which inverses the phase of the right channel to produce the out-of-speakers effect (Track 32.186)—neither center nor side panning seemed ideal for this mono track. It is also sent to the ambiance reverb, which blends it into the depth field (Track 32.187).

ES2 Rise

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

ES2 Rise (Track 32.188) appears once in the arrangement toward the end of the first break. It is untreated, but sent to the ambiance reverb. Apart from blending it with the depth field, some of this track's definition is courtesy of the reverb and its tail (Track 32.189).

Sine Rise

Inserts: MOTU *MasterWorks EQ*, MOTU *Delay*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Sine Rise (Track 32.190) only plays once toward the end of the second break along with Riser, Reversed Rise, and Bent Phase, and thus its role is somewhat redundant. A choice had to be made as to which tracks would take the spotlight: Riser and Reversed Rise or Sine Rise. Since Riser and Reversed Rise are richer, Sine Rise ended up being layered underneath them. Sine Rise first goes through the *MasterWorks EQ*, where a parametric filter [–6 dB, 4 kHz, *Q* 0.83] tunes it into the frequency spectrum (Track 32.191). The originally narrow stereo width is made wider by the MOTU *Delay*, which delays the right channel only by 80 ms (Track 32.192). It is sent to the ambiance reverb, which sends it quite deep into the depth field (Track 32.193).

Strings

Curtis Strings

Inserts: MOTU *MasterWorks EQ*.

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Curtis Strings (Track 32.194) is only treated with the *MasterWorks EQ*, where an HPF [100 Hz, 18 dB/oct] rolls off dispensable lows (Track 32.195). This track is also sent to the ambiance reverb (Track 32.196) and its level is automated between the various sections.

Strings Finale

Sends: Ambiance Reverb (Audio Ease *Altiverb*).

Strings Finale (Track 32.197) concludes the production. This track is untreated, but to add some drama, its imprinted rise in level is enforced by level automation (Track 32.198). It is also sent to the ambiance reverb (Track 32.199).

33 "Hero" (rock)

Performed by AutoZero.
Lyrics by Dan Bradley, music by AutoZero.
Dan Bradley: vocals, guitars.
Lish Lee: bass, vocals.
Lee Ray Smith: drums.
Produced by AutoZero.
Recorded at Soho Studios, London.
Engineered by Guy Katsav.
Edited by Luca Barassi.
Mixed (in Pro Tools) by Roey Izhaki.
Mastered by Mandy Parnell (www.mandyparnell.com) at Electric Mastering, London.

AutoZero is a three-piece indie band from London. "Hero" is a rock song based on classic rock instrumentation—drums, bass, rhythm/power guitars, and vocals. The song was recorded overnight at Soho Studios, London. Although offering state-of-the-art mixing facilities, Studio 1 (where "Hero" was recorded) has a very small live room—just big enough for a drum kit. This small room is evident in the drum recordings. The limited time—7 hours—in which this production was recorded is another contributor to the fact that the raw material was less than ideal. As with many other recordings made in limited time, much of the mix efforts focused on correcting rather than elevating.

Track 33.1 is a home-recorded demo of the song, involving an electronic drum kit. Track 33.2 is a mix-ready version involving rough levels and panning, but no processors or effects. Track 33.3 is my own rough mix of the pre-edited multitrack. Neither the toms nor the backing vocals are mixed in this rough mix, and the work on it stopped once I felt familiar enough with the tracks and had a firm plan for the real mix. Perhaps the greatest lesson offered by the rough mix was that some distinction had to be made between the verses and the choruses—the original recordings involved little arrangement changes between the sections. Track 33.4 is a mix-in-progress version, Track 33.5 is the final mix, while Track 33.6 is the mastered final version.

Drums

Overheads

Inserts: Sonnox *Oxford EQ*, PSP *Vintage Warmer*.

The small live room in which the drums were recorded resulted in an overheads recording that had a few flaws (Track 33.7). The small-room ambiance captured on the recording was unappealing. In addition, the individual kick, snare, and tom tracks all suffered from poor interaction with the overheads. For these reasons, the cymbals-only approach was chosen for the overheads. First, they went through the *Oxford EQ*, where a steep HPF [400 Hz, 30 db/oct] goes as high as it can to roll off all the lows (Track 33.8). In addition, a parametric filter [702 Hz, –8.37 dB, *Q* 3.08] further attenuates the snare and removes an annoying mid-range noise (Track 33.9). The EQ is followed by the *Vintage Warmer* [drive +12.9 dB, knee 41.1 percent, speed 35.8 percent, auto release, mix 78 percent], where parallel compression tightens the sound (Track 33.10). The level of the overheads was automated to rise during the last drum bar. Also, their stereo image was narrowed to give a tighter impression and to clear space on the extremes for the distorted guitars.

Kick

Kick 1 Inserts: Toontrack *EZdrummer*, Digidesign *Smack!*, Sonnox *Oxford EQ*.

Kick 1 Sends: Ambiance Reverb (UAD *DreamVerb*).

Kick 2 Inserts: Toontrack *EZdrummer*, Digidesign *Smack!*, Sonnox *Oxford EQ*.

The multitrack contained both kick-in and kick-out tracks (Tracks 33.11 and 33.12). Having auditioned these tracks while working on the rough mix, it quickly became apparent that replacing the recorded kick with triggered samples would not only provide easier material to work with, but would also yield a better sound altogether. The kick on the rough mix is already triggered. The triggering process involved importing the kick-in track into Logic and converting it to MIDI using the "Audio to Score" feature. The resultant MIDI track was then loaded on an instrument track in Pro Tools, with *EZdrummer* as a virtual instrument. The kick sound used is the default 22″ GMS Felt Beater from the Pop/Rock kit.

In the final mix, there are two kick tracks. The two never play together and the toggling between them sharpens the more powerful moments in the song, as well as adding some interest. Both kick samples are identical and both are processed using the same plugins, but the settings on these plugins are different.

Kick 1 (Track 33.13) is the more rounded, settled kick and it plays through most of the verses and during the break. It is first compressed using *Smack!* [norm mode, input 5, ratio 2:1, attack 4.1, release 0.8], which adds some punch and mass (Track 33.14). Then it goes through the *Oxford EQ*, where four parametric filters shape the sound of the kick: the first filter [83.3 Hz, 3.24 dB, *Q* 2.83] adds some thud (Track 33.15), the second filter [220 Hz, –5.26 dB, *Q* 2.57] attenuates tapping low-mids (Track 33.16), the third filter [1,003 Hz, 5.6 dB, *Q* 2.83] adds attack (Track 33.17), and the fourth filter [5 kHz,

–2.87 dB, *Q* 1.93] reduces click (Track 33.18). This kick is also sent at a low level to the ambiance reverb, partly to soften it, partly to blend it into the ambiance, partly to move it farther back in the mix, and partly to add to its timbre (Track 33.19).

Kick 2 (Track 33.20) is the more powerful kick between the two, having more snap and presence. It plays during the choruses, bridge, outro, and some verse parts. It is first compressed using *Smack!* [norm mode, input 6.4, ratio 6:1, attack 10, release 0], which adds tightness and punch (Track 33.21). Also, the distortion facility in *Smack!* [odd__+ even] was used to add some grit (Track 33.22). The compressor is followed by the *Oxford EQ*, where four parametric filters shape the sound of the kick: the first filter [110.5 Hz, 3 dB, *Q* 2.83] adds thud (Track 33.23), the second filter [260 Hz, –5.78 dB, *Q* 1.6] reduces unnecessary mids (Track 33.24), the third filter [938 Hz, 11.45 dB, *Q* 4.83] adds attack (Track 33.25), and the fourth filter [5.1 kHz, 2.4 dB, *Q* 3.9] adds click (Track 33.26).

Snare top

Inserts: PSP *Neon*, Digidesign *DigiRack EQ 3*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

Having tried gating, compression, and triggers, I came to the conclusion that equalization alone was the most suitable approach for the snare top. A healthy amount of hi-hats and ride spill on this track made the cymbals unstable once compression or gating had been applied, and a fully triggered snare sounded too mechanical. Not gating the snare meant that spill on the track limited equalization possibilities, mainly with regard to the highs, where the hi-hats and the ride roam. On the other hand, there was some benefit in the support the hi-hats got from the snare track.

Snare top (Track 33.27) first goes through *Neon*, which in linear-phase mode shapes its sound. A low-shelving filter [1.1 kHz, –5.9 dB] attenuates an excess of lows and low-mids (Track 33.28), a parametric filter [312 Hz, –4 dB, *Q* 1.8] attenuates a resonant tone (Track 33.29), and another parametric filter [2.5 kHz, –5 dB, Q 2.9] attenuates some cymbals spill without harming the snare timbre too much (Track 33.30). On the *DigiRack EQ*, a high-shelving filter [4.6 kHz, –1.8 dB] softens the snare and reduces the cymbals spill (Track 33.31). The high-shelving filter was automated to only take effect during the verses—in other sections, the snare becomes brighter and the cymbals slightly louder. The sole reason for the addition of another EQ was that automating it was more straightforward than automating the *Neon*.

Snare top is sent to the ambiance reverb, which gels it to the sound stage and shifts it slightly backward (Track 33.32). The level of this track was automated to balance specific hits that were either too loud or too quiet. The track is panned to approximately 10:30.

Snare bottom

Inserts: Digidesign *DigiRack EQ 3*, Sonnox *Oxford Dynamics*.

Sends: Snare Reverb (Audio Ease *Altiverb*, McDSP *Channel G Dynamics*).

Snare bottom has somewhat more importance in this track since snare top did not deliver a complete snare sound. Snare bottom complements snare top and

contributes both presence and crispness (Tracks 33.33 and 33.34). The raw track (Track 33.35) first goes through the *DigiRack EQ*, where an HPF [313 Hz, 24 dB/oct] rolls off much of the kick spill and the snare's body, which is already contributed by the snare top (Track 33.36); a parametric filter [1 kHz, 5.1 dB, *Q* 1] adds some presence and attack (Track 33.37), and a high-shelving filter [6 kHz, 4 dB] brightens the snare and accentuates its crispness (Track 33.38). Following the EQ, the *Oxford Dynamics* was inserted, where both the gate and the compressor are employed. The gate [threshold –19.2 dB, range –40 dB, attack 0.01 ms, hold 30 ms, release 11.2 ms] removes spill, notably the snare rattle caused by kick hits (Track 33.39). The compressor [threshold –18 dB, ratio 2:1, soft-knee 40-dB span, attack 5.2 ms, release 130 ms] adds some weight and density (Track 33.40). Snare bottom is panned to the same position as snare top.

Snare bottom is sent to *Altiverb*, and the IR used is the "Plate Short" from the EMT 140 (Wendy Carlos) category (Track 33.41). The only modification on the reverb is a 9.2-dB boost on the equalizer's treble (Track 33.42), which gives the reverb more shine. The reverb is gated by the *Channel G Dynamics*, with its key input set to the dry snare (Track 33.43). This gated reverb was added as an effect, but also in order to send the snare farther back.

Snare triggers

Main Track Inserts: Digidesign *Smack!*, Digidesign *DigiRack EQ 3*.

Support Track Inserts: Digidesign *DigiRack EQ 3*, McDSP *Compressor Bank CB2*.

A triggered snare track was generated using Trillium Lane Labs *Drum Rehab*. From its integrated library, the sample chosen was the DW1 snare from the "Perfect Drums" collection. Although ultimately the snare sound is based on the recorded snare-top track, the triggered snare is still mixed, but only as a layer underneath the recorded snare. The triggered snare is mixed in all sections of the song apart from the verses. Its main role is to add brightness and definition to the recorded snare sound (Tracks 33.44 and 33.45). Also, since the velocities of the triggered hits are nearly identical, they balance out to some extent the overall snare hits.

The raw triggered snare (Track 33.46) had a decay that was too long and yielded a fluffy overall snare sound. So it is first treated by *Smack!* [norm mode, input 6.3, ratio 3:1, attack 5.8, release 4.7], which adds punch using moderate attack but also attenuates the decay using moderate release (Track 33.47). The compressor is followed by the *DigiRack EQ*, where a combination of an HPF [262 Hz, 12 dB/oct] and a deep dip on a parametric filter [427 Hz, –14 dB, *Q* 1] removes all the snare body and tone (Tracks 33.48 and 33.49); a high-shelving filter [6 kHz, 1.7 dB] adds a touch of highs as part of tonality shaping (Track 33.50).

From the triggered snare track, a few hits were copied onto a new track. These provide extra support during the final sections of the song, mainly boosting the overall snare level but also causing a subliminal change to the snare tonality. These few hits appear right before the break, right after it, halfway through the outro, and at the very end. These hits (Track 33.51) are first treated with the *DigiRack EQ*, where an HPF [476 Hz, 6 dB/oct] rolls off much of the lows (Track 33.52); a high-shelving filter [4.61 kHz,

–3.4 dB] attenuates the highs so these hits are not be over-present (Track 33.53). The equalizer is followed by the *CB2* [threshold –30.6 dB, comp 7.6, knee 15, attack 70 ms, release 273 ms], which adds attack (Track 33.54).

Hi-hats

Inserts: Digidesign *DigiRack EQ 3*.

Having the hi-hats sufficiently present on the overheads and snare top meant that there was little point mixing the hi-hats track. Still, the hats could use a little push during the break, so this was the only place where this track was mixed (Track 33.55). The hi-hats only go through the *DigiRack EQ*, where an HPF [668 Hz, 18 dB/oct] removes kick spill and low-mids that contributed little to the sound (Track 33.56), and an LPF [6.8 kHz, 6 dB/oct] eases brittle highs (Track 33.57).

Ride

As there was sufficient presence on the overheads, this track was omitted.

Toms

First Layer Inserts: McDSP *Compressor Bank CB2*, Digidesign *DigiRack EQ 3*.

First Layer Sends: Ambiance Reverb (UAD *DreamVerb*).

Second Layer Inserts: McDSP *FilterBank E2*, Digidesign *Smack!*

Like the kick, the toms in the mix are triggered. But at one specific section, just before the break (Tracks 33.58 and 33.59), the recorded toms did an exceptional job. The tom sound in this section is the outcome of two layers (Tracks 33.60–33.62); both involve the original tom recording (Track 33.63), but each is processed differently. It is the level balance between these two layers that dictated the final toms sound.

The first layer first goes through the *CB2* [threshold –37.3 dB, comp 3.4, knee 13.7, attack 70 ms, release 40.2 ms], which was employed to condense the natural-sounding dynamics (Track 33.64). It then goes through the *DigiRack EQ*, where a parametric filter [3.5 kHz, 5 dB, *Q* 0.81] adds attack and highs that let the toms push more through the mix (Track 33.65). This layer is also sent to the ambiance reverb—although the reverb does not handle the toms very well, once mixed it adds a rich sense of warm space and a touch of chaos, which works well (Track 33.66).

The second tom layer first goes through the *FilterBank E2*, where a low-shelving filter [968 Hz, –12 dB] drys the lows and low-mids, leaving only the presence and attack portions of the toms (Track 33.67). The EQ is followed by *Smack!* [input 7, ratio 6:1, attack 8.5, release 6.1], which accents the attack and sharpens the tom's dynamics (Track 33.68).

Tom triggers

Tom 1 Inserts: Digidesign *Smack!*, McDSP *FilterBank E6*.

Tom 1 Sends: Ambiance Reverb (UAD *DreamVerb*).

Floor Tom Inserts: Digidesign *Smack!*, McDSP *FilterBank E6*.

Floor Tom Sends: Ambiance Reverb (UAD *DreamVerb*).

With the exception of the short section above, all other tom hits are triggered. This allows the toms to project better and provides ultimate control over their sound.

Tom 1 is used several times during the song. The tom hit (Track 33.69) was extracted from isolated hits recorded in the studio after the drums were tracked. These hits were pasted into the session manually using the tab-to-transient feature. They first go through *Smack!* [input 8.3, ratio 6:1, attack 3.9, release 0], where a moderate attack adds some punch while zero release lengthens the decay (Track 33.70). The distortion facility on *Smack!* [odd] was also utilized to add a touch of grit (Track 33.71). The compressor is followed by the *FilterBank E6*, where three bands are used: an HPF [135 dB, 12 dB/oct] sets the low frequency limit of the tom, essentially preventing it from overloading the lows, including those of the reverb it is sent to (Track 33.72); a parametric filter [648 Hz, –2.6 dB, *Q* 0.8] shapes the tonality of the tom while removing dispensable mids (Track 33.73); and another parametric filter [4.35 kHz, 2 dB, *Q* 1] accents the attack (Track 33.74). Tom 1 is sent to the ambiance reverb, mainly in order to detach it from the front of the mix (Track 33.75).

The floor tom (Track 33.76) only plays once just before the bridge. Like tom 1, it was also extracted from the isolated hits recorded in the studio. It is compressed quite similarly to tom 1 by *Smack!* [input 8.3, ratio 6:1, attack 6, release 0], which adds attack and lengthens the decay (Track 33.77). On *FilterBank E6*, which follows, a single parametric filter [4.9 kHz, 6.9 dB, *Q* 0.9] accents the attack and definition of the tom (Track 33.78). The floor tom is sent to the ambiance reverb, intentionally creating some distant thunderous thud (Track 33.79).

Tambourine

Inserts: Toontrack *EZdrummer*, Digidesign *DigiRack Dyn 3*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

Break Tambourine Inserts: Toontrack *EZdrummer*, Sonnox *Oxford Reverb*.

The tambourine in "Hero" was never recorded—instead, it was generated using *EZdrummer* during mixdown. It plays through most of the song except the verses, and was added mainly to sharpen the contrast between the verses and the choruses. It was mixed as a support instrument, meaning it is not clearly defined (Tracks 33.80 and 33.81). The untreated tambourine (Track 33.82) was compressed by the *DigiRack Dyn 3* [threshold –12.6 dB, ratio 5.3:1, soft-knee 24.7 dB span, attack 40.6 µs, release 5.7 ms], which balances out the loud strokes with the quiet shakes (Track 33.83). It is sent to the ambiance reverb, which sends it far back into the depth field (Track 33.84). Before the verses, the tambourine is panned left opposite the ride, which plays along with it. From the bridge onward, the tambourine plays with the hi-hats, which appear left. To prevent the hi-hats from masking the tambourine and in order to create a more balanced stereo panorama, the panning of the tambourine was automated to shift to the right from the bridge onward.

During the break, the tambourine's nature alters, and because of this a new track was created. The only plugin inserted on this track is the *Oxford Reverb* [67 percent wet], where a modified version of the "Concert Room" preset adds a long and shiny tail (Tracks 33.85 and 33.86). This track, which is panned center, was automated in level to rise toward the end of the break, building up to the explosive outro.

Bass

Main bass

Inserts: PSP *Vintage Warmer*, Sonnox *Oxford Dynamics*, Digidesign *DigiRack EQ 3*, SoundHack *+Chebyshev*.

As I was growing to like the bass sound in the rough mix, most of the plugins and their settings were transported from the rough mix to the final mix. The raw bass (Track 33.87) first goes through the *Vintage Warmer* [mix 68.4 percent, drive 11.6 dB, knee 31.9 percent, speed 91 percent, auto release], which has two roles: it condenses the bass dynamics using parallel compression and adds warm saturation. This saturation yields added harmonics that sharpen the bass definition (Track 33.88). Following the *Vintage Warmer* is the *Oxford Dynamics*, but the only facility used on it is Warmth [89 percent], which adds some clarity while also enlivening the sound (Track 33.89). Then comes the *DigiRack EQ*, where three bands are utilized: an HPF [72.4 Hz, 24 dB/oct] was employed to remove low frequencies that jumped with very low notes, making the bass tonality inconsistent (Track 33.90); a parametric filter [206.5 Hz, –1.6 dB, *Q* 1] was added late in the mixing stage to ease the mids on the bass as part of frequency tuning (Track 33.91); and another parametric filter [636 Hz, 3.8 dB, *Q* 1] simply shapes the sound of the bass by accenting low-order harmonics (Track 33.92). Following the EQ comes the *+Chebyshev*, which adds even more distortion and a very distinct size (Track 33.93).

Break bass

Inserts: PSP *Vintage Warmer*, Digidesign *DigiRack EQ 3*, SoundHack *+Chebyshev*, Sonnox *Transient Modulator*.

During the second part of the break, the bass plays an important solo role, which builds up to the outro. This called for a deviation from its normal sound. So the bass during the break was moved to a new track, which was processed differently from the main bass. This track was first treated with the *Vintage Warmer* [mix 100 percent, drive 20.9 dB, knee 31.9 percent, speed 91 percent, auto release], which tightens the sound while also adding some warmth (Tracks 33.94 and 33.95). Then comes the *DigiRack EQ*, where an HPF [124 dB, 18 dB/oct] intentionally dries the lows from the bass, with the idea that these will return with the outro (Track 33.96); a parametric filter [340 Hz, –7.5 dB, *Q* 2.24] attenuates disturbing mids (Track 33.97); another parametric filter [497 Hz, 8.8 dB, *Q* 10] opens up the flat sound and adds some resonance as an effect (Track 33.98); and yet another parametric filter [2.46 kHz, 10.8 dB, *Q* 4.06] draws some presence (Track 33.99). Next comes the *+Chebyshev*, which adds an extra dimension to the sound (Track 33.100). Last comes the *Transient Modulator* [ratio 0.25, overdrive 90 percent], which accentuates attack and adds a subtle amount of harmonic content (Track 33.101).

Bass FX

Inserts: Digidesign *DigiRack EQ 3*, SansAmp *PSA-1*, McDSP *Chrome Stack*, Digidesign *DigiRack EQ 3*.

As part of early tryouts, I experimented with the idea of having a heavily distorted bass, an idea that did not quite work well in the context of this song (Track 33.102). Yet I had the distorted layer (a duplicated track of the bass going through the following plugins) and it worked quite well during the powerful outro of the song—adding extra power to the bass and to the overall mix (Tracks 33.103 and 33.104). The first plugin in the chain is the *DigiRack EQ*, which acts as a tuner for the subsequent *PSA-1* (Tracks 33.105–33.107). Two bands are operational on the EQ: an HPF [194.2 Hz, 24 dB/oct] and a high-shelving filter [8.81 kHz, –6 dB]. The *PSA-1* is followed by the *Chrome Stack*, which contributes the main effect (Track 33.108). As part of the final frequency tuning, another *DigiRack EQ* was inserted, where a parametric filter [829.6 Hz, –2.9 dB, Q 1] pulls some mids (Track 33.109).

Rhythm guitar

Three takes of the rhythm guitar were recorded, each using three sources: a DI recording, an on-the-grill Shure *SM57* and a distant Neumann *U87* (Tracks 33.110–33.112). Out of the three, the DI and *SM57* were used in the mix, while the distance captured on the *U87* recording proved too limiting and wasn't used.

The rhythm guitar plays a more dominant role during the intro and break compared with the rest of the song, so a dedicated track was allocated for these two sections.

Intro/break rhythm guitar

Inserts: Digidesign *DigiRack Dyn 3*, SansAmp *PSA-1*, Digidesign *DigiRack EQ 3*, Sonnox *Oxford Reverb*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

The intro/break guitar, a DI recording (Track 33.113), first goes through the *DigiRack Dyn 3*, where the compressor [threshold –26.4 dB, ratio 9.7:1, soft-knee 30-dB span, attack 1 ms, release 15.1 ms] evens out level fluctuations (Track 33.114). It is followed by the *PSA-1*, which adds some color to the pale DI sound while also adding some crunch (Track 33.115). Next comes the *DigiRack EQ*, where a low-shelving filter [230.6 Hz, –7.2 dB, *Q* 1] attenuates an excess of lows (Track 33.116), and then a wide parametric filter [2 kHz, 7.4 dB, *Q* 1] adds essential highs to fabricate the final guitar sound (Track 33.117). The *Oxford Reverb* was inserted in series to add some stereo size to the mono recording. The reverb was programmed to only output early reflections with very little (8 percent) of the wet signal mixed (Track 33.118). In addition, this track was sent to the ambiance reverb in order to place it in the mix space (Track 33.119).

The intro guitar plays for four bars alone before other instruments are introduced—an introduction that in my opinion did not have sufficient power. So the level of the intro

guitar was set intentionally at a level low enough to allow some level impact once other instruments are introduced (compare Tracks 33.120 and 33.121).

Main rhythm guitar

Inserts: Digidesign *DigiRack EQ 3*, Digidesign *DigiRack Dyn 3*.

Sends: Guitar Chorus (UAD *Delay Modulator DM-1*), *Ambiance Reverb* (UAD *DreamVerb*).

By the time all the tracks had been brought up and coarsely mixed, there was very little space for the rhythm guitar. Its definition is in the high-mids—an area already occupied by virtually all other instruments but the bass. Luckily, a rhythm guitar need not be in the limelight as it plays a rather supportive role in the mix.

The raw track (Track 33.122), a DI recording, went through aggressive equalization using the *DigiRack EQ*. An HPF [648.7 Hz, 18 dB/oct] rolls off a good portion of the lows and low-mids, which contributed little to the guitar sound yet cluttered a valuable range of frequencies (Track 33.123). A parametric filter [1.1 kHz, –8.3 dB, *Q* 1] in combination with another parametric filter [1.83 kHz, –9.5 dB, *Q* 3.79] tunes the guitar into the busy high-mids (Tracks 33.124 and 33.125). Then a high-shelving filter [6 kHz, –2.7 dB] treats exaggerated highs (Track 33.126). The EQ is followed by the *DigiRack Dyn*, where a compressor evens out level fluctuations (Track 33.127). The guitar is sent to the *DM-1*, which by adding chorus effect sends the guitar back in the depth field and also creates a wide curtain of sound (Track 33.128). The guitar is also sent to the ambiance reverb, which gels it into the sound stage (Track 33.129).

Support rhythm guitar

Inserts: McDSP *FilterBank E6*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

The main rhythm guitar was panned to the right, yet, despite the stereo effects added, something was still missing on the left. To combat this, a different take—an *SM57* recording—was mixed panned left. Combined with the main rhythm guitar, this additional track creates an even wider and richer impression (Tracks 33.130 and 33.131). The raw track (Track 33.132) was only equalized using the *FilterBank E6*, where an HPF [1.23 kHz, 6 dB/oct] rolls off dispensable lows and low-mids (Track 33.133). It is sent to the ambiance reverb, which places it on the same line as the main rhythm guitar (Track 33.134).

Distorted guitars

When it came to mixing the distorted guitars, a few things had to be considered. First, during the verses, there are intensity changes between sections that involve the vocal and those involving the lead guitar. To support these changes, the distorted guitars are split into two pairs: one pair, which I term "curtain guitars," play during the relaxed sections of the verses; the other pair, simply termed "power guitars," play during all other sections. The second consideration was that the power guitars in the choruses, bridge, and outro should stand out more than those during the verses.

The distorted guitars were recorded using the same DI, *SM57*, and *U87* setup as the rhythm guitar, and here again the *U87* was not used. The DI recordings that were used were processed by amp simulators—a process that allowed ultimate control over their final sound.

Curtain guitars

Guitar 1 Inserts: UAD *Preflex*.

Guitar 2 Inserts: Digidesign *DigiRack Mod Delay II*, Digidesign *DigiRack EQ 3*.

Sends: Ambiance Reverb (UAD *DreamVerb*).

Curtain guitar 1 is a DI recording (Track 33.135) processed by the UAD *Preflex* with a modified version of the "Foxy Gravy" preset (Track 33.136). This track was panned hard-left.

Curtain guitar 2 is an *SM57* recording of the same take as guitar 1. It is panned hard-right. The raw track (Track 33.137) goes through the *DigiRack Mod Delay* [mix 50 percent, 24-ms delay, no feedback or modulation] for the sole purpose of sending it backward without adding stereo width (Track 33.138). On the succeeding *DigiRack EQ*, a low-shelving filter [175.5 Hz, –8 dB] attenuates an excess of lows (Track 33.139).

Both tracks are routed to an audio group, which is sent to the ambiance reverb. The reverb shifts the guitars farther back in the mix and also detaches them from the extremes (Track 33.140). This creates a nice contrast between the curtain guitars and the power guitars—the curtain guitars appear farther back and more centered; the power guitars appear farther in front and their image is bound to the extremes. The level of the curtain guitars was automated so they only appear in the relaxed section of the verses.

Power guitars

Guitar 1 Inserts: SansAmp *PSA-1*, SansAmp *PSA-1*, Digidesign *DigiRack Mod Delay II*.

Guitar 2 Inserts: Digidesign *DigiRack EQ 3*, Digidesign *DigiRack Mod Delay II*.

Power guitar 1 is a DI recording (Track 33.141). The first plugin in the inserts chain is the *PSA-1*, which is bypassed during the verses (Track 33.142). It is followed by another *PSA-1*, which operates throughout (Track 33.143). While the second *PSA-1* plays alone during the verse, in all other sections it is combined with the first *PSA-1*—a combination that creates a richer, brighter, and more powerful sound (Track 33.144). Next comes the *Mod Delay* [36-ms delay on right channel only], which, by applying the Haas trick, turns the mono track into stereo while adding a noticeable size (Track 33.145).

Power guitar 2 is an *SM57* recording of a different take than the one used for power guitar 1. The raw recording (Track 33.146) goes through the *DigiRack EQ*, where a low-shelving filter [142.4 Hz, –4.4 dB] attenuates overemphasized lows (Track 33.147) and a parametric filter [2.5 kHz, 2 dB, *Q* 1] adds a touch of presence (Track 33.148). The guitar then goes through the *Mod Delay* [57.23-ms delay on the left channel only], which turns the mono track into a wide stereo one; modulation on the delay [depth 3 percent, rate 1.26 Hz] adds further dimension and size (Track 33.149).

Distorted guitar group

Inserts: Sonnox *Oxford EQ*.

All distorted guitar tracks were routed to an audio group. On the group track, the *Oxford EQ* is responsible for tuning the guitars to the frequency spectrum. On the EQ, four bands are employed: an HPF [288.2 Hz, 18 dB/oct] sets the low-frequency limit of the guitars (Tracks 33.150 and 33.151); a parametric filter [600 Hz, –2.75 dB, *Q* 6.25] attenuates some mids that the guitars could live without and other instruments could use (Track 33.152); for the same reason, another parametric filter [3.7 kHz, –2.75, *Q* 2.22] clears some highs (Track 33.153); and then a high-shelving filter [9.4 kHz, –6.28 dB] pulls some harsh highs (Track 33.154).

Lead guitar

Lead guitar

Inserts: Digidesign *DigiRack Dyn 3*, Digidesign *DigiRack EQ 3*, PSP *Nitro*.

The selected source for the lead guitar was the *U87* recording (Track 33.155). It is first compressed using the *DigiRack Dyn* [threshold –48.9 dB, ratio 3:1, hard-knee, attack 272.7 μs, release 13.6 ms] in order to balance out level fluctuations and also to increase sustain (Track 33.156). The compressor is followed by the *DigiRack EQ*, where a steep HPF [448.6 Hz, 24 dB/oct] rolls off dispensable lows and low-mids (Track 33.157); a parametric filter [2.09 kHz, 4.5 dB, *Q* 1.38] adds some flesh (Track 33.158); and another parametric filter [5.96 kHz, –4.1 dB, *Q* 3.23] eases some harshness, resulting in an overall smoother, more rounded sound (Track 33.159). The sound of the lead guitar is mostly the outcome of the *Nitro* and its "No-Fi" preset—essentially a combination of bit reduction and sample rate reduction followed by an LPF (Tracks 33.160 and 33.161). The lead guitar is sent to the ambiance reverb, mainly in order to shift it backward in the depth field and to blend it into the sound stage (Track 33.162). The lead track level was automated at various sections of the song, for example, to rise during the intro.

Lead guitar FX

Inserts: Digidesign *DigiRack EQ 3*, Digidesign *Smack!*, PSP *84*.

Another element that was transported from the rough mix was the lead guitar effect. To prevent correlation to the lead track, the original lead track (Track 33.163) was duplicated and processed as follows: first, it is equalized by the *DigiRack EQ*, where an HPF [623.3 Hz, 6 dB/oct] rolls off dispensable lows (Track 33.164) and a high-shelving filter [6 kHz, 3.6 dB] adds some spark (Track 33.165). *Smack!* [norm mode, input 5, ratio 4:1, attack 2.9, release 5] is then in charge of balancing level variations and lengthening the sustain (Track 33.166). The wah-wah-like effect is the outcome of a modified "Sarawak" preset on the PSP *84* (Track 33.167). This track is also sent to the ambiance reverb (Track 33.168).

Vocals

Lead vocal

Inserts: Pultec *EQH-2*, McDSP *MC2000 MC2*, SoundHack *+Chebyshev*, Sonnox *Oxford Dynamics*.

Sends: Ambiance Reverb (UAD *DreamVerb*), *Vocal Reverb* (UAD *DreamVerb*).

There was a conceptual issue with mixing the lead vocal. On the one hand, "Hero" has its share of aggressiveness; on the other, it has its pop characteristics and involves a catchy melody. So the main question was: Should the lead vocal be sweet and beautiful, or should it be bad and dirty? It ended up somewhere in between.

The first plugin in the inserts chain, the *EQH-2*, was employed to shape the basic tonality of the vocal. It removes dispensable lows [100 Hz, attenuation 4.9] and adds a touch of spark [10 kHz, boost 1.5] (Tracks 33.169–33.171).

Following the EQ is the *MC2000*—a multiband compressor. I tried two different compressors (*Smack!* and the *Oxford Dynamics*) prior to trying the *MC2000*. *Smack!* gave an obvious character that I found was too much in the context of the song; *Oxford Dynamics* did a fair job, but not a perfect one. Dan's voice fluctuates noticeably during some sections of this song. Whatever settings I tried, the level of the vocal never seemed stable throughout the song. The fundamental problem was that the frequency content of the recording changes quite noticeably in relation to the pitch sung. Lower-pitch notes seemed to disappear in the mix—making the vocal come and go. This issue is not uncommon, but on this production it seemed extremely tricky. For this reason, I picked the *MC2000*. Perhaps the most important setting on the compressor was the crossover frequency, which ended up at 1.6 kHz. Both bands are heavily compressed, with the threshold of the low band set to –24 dB and the threshold of the high band to –44 dB. In addition to their standard use, the independent gain controls of each band were used to further shape the tonality of the vocal (Tracks 33.172–33.174). The *MC2000* added quite some character to the vocal, which I liked.

Following the *MC2000*, the vocal sounded balanced but still did not stand out as much as it should have. This was mostly the outcome of missing vitality on the mids. The *+Chebyshev* did a perfect job of filling this missing range. It also added some aggression in the form of distortion that at points becomes rather obvious—an effect that I thought was for the better (Track 33.175). The last plugin in the chain is the *Oxford Dynamics* [threshold –12 dB, ratio 2:1, soft-knee 40 dB span, attack 0.52 ms, release 5 ms], which acts as a catch-all compressor to perfect the balancing of the distorted voice (Track 33.176).

The lead vocal is sent both to the ambiance reverb and to a dedicated vocal reverb. The ambiance reverb blends the vocal into the sound stage, while the vocal reverb—a slightly modified "PLATE 140 Vocal Plate" preset—adds a rich effect, life, and size (Tracks 33.177–33.179). The send level to the vocal reverb was automated to rise for the duration of a few sustained vowels (Track 33.180).

Vocal FX

The lead vocal involves a few effects. First, there is a narrowed ping-pong delay after the first chorus. This delay was not produced by a plugin, but by crossed bus routing in Pro Tools—left bus sent to right bus, right bus sent to left bus. On each bus there is 830 ms of delay, which, despite not being perfectly synced to the tempo, sat quite right. In addition, the echoes were treated by the *DigiRack EQ 3*, which was inserted on each bus with a combination of an LPF and a parametric filter that creates a resonant filter. The send level on these buses was automated so only three echoes are produced—a fourth echo would clash with the verse vocal. The echoes are blended into the depth field by the ambiance reverb to which they are sent (Track 33.181).

The ping-pong delay closing the first chorus fills some space between the chorus and the second verse. The second chorus is followed by an additional vocal phrase and then the lead guitar immediately takes the stage, so there was little place for the same ping-pong effect. The only processing on the vocal phrase after the second chorus is the SansAmp *PSA-1* (Track 33.182).

Halfway through the bridge, the phrase "deep within us" is also processed, this time using four plugins. The first plugin, which is the main effect contributor, is the PSP *Master Q*. The settings involve an HPF at 96 Hz, an LPF at 3.5 kHz, and a mad boost of 24 dB at 981 Hz. The actual effect is also the outcome of the hard saturation on the saturation facility (Tracks 33.183 and 33.184). The *Master Q* is followed by the Lexicon PSP *42*, which adds a 172-ms delay with feedback (Track 33.185). Then comes the Digidesign *D-Verb* set to large hall and 19 percent wet (Track 33.186). Finally, the McDSP *Chrome Tremolo* adds a subtle tremolo effect (Track 33.187). This specific phrase is sent to the ambiance reverb (Track 33.188).

The "deep within us" phrase shortly after repeats. This second instance could have been processed in the same manner, but to create some variation and contrast a different effect was used: the lead vocal sends to the ambiance, and vocal reverbs are muted, causing forward shifting of the voice (Track 33.189).

Fake second voice

Inserts: Lexicon PSP *42*, Digidesign *DigiRack EQ 3*.

Sends: Vocal Reverb (UAD *DreamVerb*).

One of the last overdubs to be recorded, at around 6 a.m., was the second voice. By that point, Dan's vocal cords were a few takes away from snapping. Second voice harmony was recorded for all but the verses, but the performance during the choruses and the outro was pretty unconvincing. As a replacement, the lead vocal track was duplicated, trimmed in the right places, and processed with ADT settings on the PSP *42* [15-ms modulated delay, no feedback] (Tracks 33.190–33.192). Following the PSP *42* comes the *DigiRack EQ*, where an HPF [239.4 Hz, 18 dB/oct] rolls off dispensable lows (Track 33.193) and a parametric filter [1.79 kHz, 7.6 dB, *Q* 2.81] boosts some presence (Track 33.194). The track is sent to the vocal reverb, which sends it backward in the depth field and also correlates it to the lead vocal (Track 33.195). To distinguish it from the lead vocal, it was panned left to around

9:00. As part of a gradual buildup, the second voice is only introduced during the second verse.

Real second voice

Inserts: PSP *84*, Sonnox *Oxford Dynamics*.

Sends: Vocal Reverb (UAD *DreamVerb*).

The only section in which the recorded second voice is featured is the bridge. The original recording (Track 33.196) was first processed with the PSP *84*, but no delay was actually applied. Instead, the filter, drive, and plate reverb were used to create a unique effect (Track 33.197). Following the PSP *84* is the *Oxford Dynamics*, where a compressor [threshold –26.5, ratio 3.48:1, soft knee 30-dB span, attack 0.52 ms, release 5 ms] contains level fluctuations (Track 33.198). The vocal is sent to the vocal reverb (Tracks 33.199 and 33.200).

Ahh BV

Inserts: Digidesign *DigiRack EQ 3*, McDSP *Compressor Bank CB1*, PSP *Nitro*.

Audio Group Inserts: Digidesign *DigiRack Mod Delay II*, McDSP *FilterBank E2*, SoundHack *+Chebyshev*.

Two backing vocals tracks were recorded for the bridge (Track 33.201) and, like the second voice performance, they are not the greatest. A conscious decision was made to bury these in the mix. The two tracks were processed identically and then routed to an audio group for additional processing. The first plugin is the *DigiRack EQ*, where an HPF [260.6 Hz, 12 dB/oct] rolls off dispensable lows (Track 33.202). It is followed by the CB1 [threshold –40 dB, comp 3.2, knee 4.7, attack 5 µs, release 204 ms], which flattens the levels (Track 33.203). The succeeding PSP *Nitro* was employed to create an autopan effect (Track 33.204).

On the audio group, the *Mod Delay* [23.75-ms delay of the right channel only, 44 percent depth and 0.74-Hz modulation] produces the Haas trick (Track 33.205). A low-shelving filter [1.7 kHz, –5.1 dB] on the *FilterBank* tunes the vocals into the mix (Track 33.206). Then the *+Chebyshev* adds some edge to what was a somewhat numb sound (Track 33.207).

Break BV

Inserts: Digidesign *DigiRack EQ 3*, Digidesign *DigiRack Dyn 3*.

Audio Group Inserts: PSP *84*.

Three backing vocal tracks were recorded from the break onward, and these also conclude the production. Two of the tracks were performed by Dan and one by Lish. Lish's voice was mixed to be the leading voice out of the three, panned a notch to the left; Dan's vocals were each panned nearly fully to opposite extremes. The three voices were first equalized using the *DigiRack EQ 3*, then compressed using the *Dyn 3*, before being routed to a bus for additional processing. Perhaps the most interesting

processing is the EQ on Lish's voice, where an HPF [505.6 Hz, 12 dB/oct] rolls off an excess of lows (Tracks 33.208 and 33.209), a parametric filter [2.93 kHz, –4 dB, *Q* 1] softens some edge (Track 33.210), and a high-shelving filter [3.67 kHz, 7.4 dB] boosts a generous quality of air and definition (Track 33.211).

The effect applied on the audio group is a slightly modified version of the PSP *84*'s "Gasherbrum Four" preset, set to 50 percent wet (Tracks 33.212 and 33.213). The group on which the PSP *84* was loaded is sent to the ambiance reverb (Track 33.214).

APPENDIX A

The science of bouncing

Bouncing is the process of recording the mix or a submix (i.e., a stem). This lets us free up the resources used for the bounced mix. For example, after mixing eight tracks of drums in an analog studio, we can bounce the drum mix onto a stereo track; then play the stereo track instead of the original raw tracks; and use any compressors, equalizers, reverbs, or other tools that were used for the drum mix.

With a console-based setup (Figure A.1), this can be done to free up channels on the desk; for example, in a situation where a digital multitrack recorder is used and there are more tracks available on the recorder than channels on the desk. Sometimes we want more than one track to be processed using the same piece of outboard gear. For example, we might want to compress both the kick and the snare using the same mono compressor; bouncing the compressed kick frees up the compressor so that it can be used for the snare. Console-bouncing usually involves routing the bounced tracks to a group, which is then sent to the associated track on the multitrack recorder. Whenever bouncing is done on a desk, it is worth remembering to fill in a recall sheet or save the mix using a recall system, so that it can be recalled later if alterations are required.
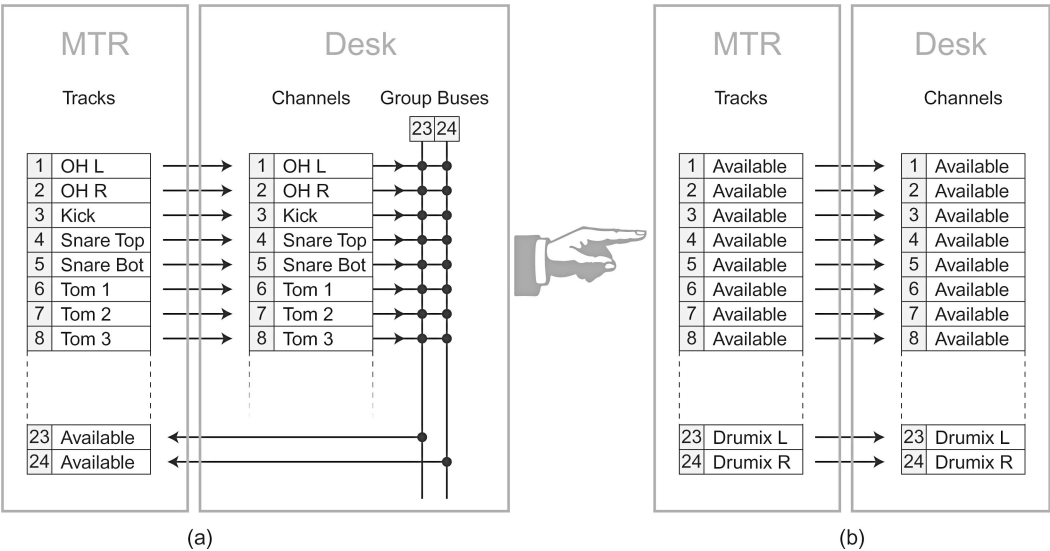
Figure A.1 Before (a) and after (b) bouncing on a console-based setup.
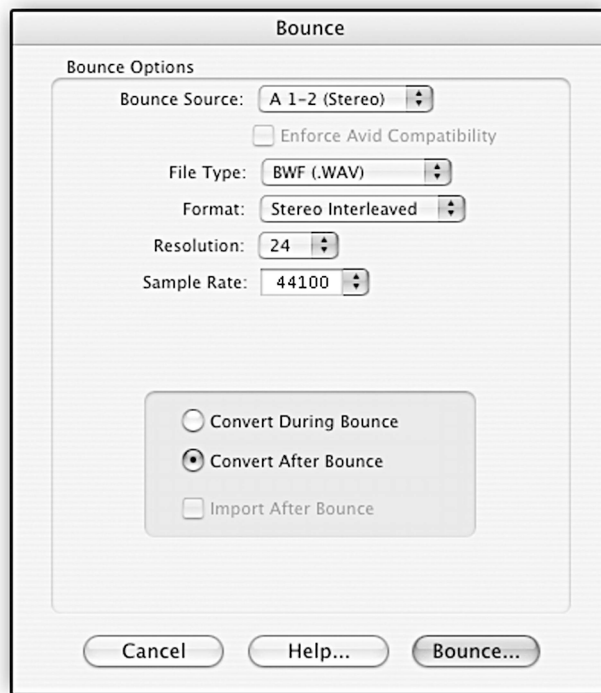
Figure A.2 Pro Tools' bounce window.

Software bouncing is usually done in order to free up CPU power so that additional plugins can be used. This is done using a dedicated command that brings up an options window (Figure A.2). In most situations, what we have to select is playback range, and solo the tracks we would like to bounce (or mute those we do not want to bounce).

The various applications provide different options in their bounce windows. Here are some options worth considering:

- **File type**—WAV files are usually the best option as these are widely supported on both PC and Mac platforms. WAV files can be timestamped (in which case they are referred to as BWFs—broadcast wave files), which can be useful when transferring projects between different applications.
- **Bit rate**—most applications' internal architecture uses a 32-bit float, which is roughly equivalent in quality to 24-bit integer (more on this in Chapter 10). Therefore, both 32-bit float and 24-bit integer are recommended. If a choice exists between the two, the native option of 32-bit float is recommended, although it will result in a bigger file. Also, 32-bit float files might not be supported at some mastering suites. Selecting a 16-bit integer might be useful when disk space is an issue, but is very likely to result in the addition of digital distortion or the dithering noise that rectifies it.
- **Sample rate**—this should always be identical to the project's sample rate.
- **File format**—the two important options are either multiple-mono or stereo-interleaved. Multiple-mono means that a stereo file is stored as two mono files

with an extension of .L or .R (e.g., *Drumix.L.wav* and *Drumix.R.wav*). Stereo-interleaved means that a stereo file is stored as a single stereo file (e.g., *Drumix.wav*). Internally, every application separates a stereo-interleaved file into its discrete left and right channels, which results in a tiny processing overhead, but makes any additional processing much faster and more manageable. Pro Tools, for instance, instead of doing this conversion in real time every time a stereo-interleaved file is played back, does it once (offline) when the files are imported into the session; this saves the tiny real-time processing overhead, but results in file redundancy that doubles up disk space. It might seem logical to bounce using the multiple-mono format, but doing so involves two issues. First, it results in more file clutter; second, not all applications (some audio editors, for example) support this format. The general recommendation here is to use the stereo-interleaved format, unless it results in redundancy, as in the case of Pro Tools.

- **Real-time (online)/offline**—with real-time bouncing, the submix is saved to a disk while being played. If we bounce a 6-minute selection, the bouncing process will take 6 minutes. While being the longer option of the two, there are a few advantages in real-time bouncing. First, it is usually less prone to timing errors that might occur when we bounce offline. The fact that we listen to what we bounce provides a form of quality control—it can reveal, for instance, clicks that were not audible before the bounced tracks were soloed. It also assures us that we are bouncing exactly the material we want (perhaps we forgot to mute a reverb return that only appears late in the song). None of this is the case with offline bouncing, where we cannot hear the result until the process is completed. As offline bouncing allocates all available processing resources to the bouncing process, this option is faster than real-time bouncing, which only uses the processing power needed for smooth playback. As a rule, in audio engineering we listen to what we commit. Therefore, real-time bouncing is recommended.

One exception to these recommendations involves the bouncing of the final mix. In this case, we usually bounce using the destination format. For example, if the file is going to be burnt onto an audio CD, it would normally be bounced using 16-bit, 44.1 kHz, stereo interleaved. But if the final mix will be mastered, **it is always better to leave any bit or sample rate conversions to the mastering stage**, since mastering engineers have better tools for such conversions. A 24-bit, 88.2 kHz project is better bounced using the same bit and sample rates, and the standard stereo-interleaved format.

Bouncing issues

Bouncing is not as straightforward as it may seem. One very common mistake when bouncing that results in quality penalty is bouncing **without checking the levels first**. In other words—using the existing mix levels rather than optimum levels. For example, if we bounce a specific pad track that on a pre-fader meter hits –5 dB and has its fader set to –15 dB, we end up with a bounced version that has its peak at –20 dB. These 20 dB of unused dynamic range result in both the analog and digital domains in smaller signal-to-noise ratio (SNR), which impairs the quality of the signal. Another issue in these circumstances is that the bounced-track fader should be at 0 dB in order for the pad to match the original mixing levels. This means that the fader can only go up by its inherent extra gain (say, 10 dB), whereas before it could go up by 25 dB (15 + 10 dB extra gain). You might think that in the digital domain we can normalize the bounced

track, but, as explained in Chapter 10, **normalization is a downgrading process** in which the SNR remains the same, but either distortion or dither noise are introduced. An opposite possible scenario involves bounced signal that is too high in level, in which case clipping distortion could be introduced. As a rule of thumb, we always want our recorded material to be at optimum levels—bounced tracks are no different.

Bouncing should be carried out at optimum levels, not mix levels.

It should be clear by this point that, prior to bouncing, we must observe the level of the bounced material and, if required, adjust it accordingly. As normally the bounced signal is that on the master bus, the master bus meters are what we look at. This helps us in determining the required gain change that will bring about the optimum level. The optimum level is the highest level possible without clipping or distorting. In the analog domain, signals can be pushed beyond the 0 VU; therefore, the definition of optimum level is not strict—but is usually accepted as being the level above which undesirable effects become apparent. In the digital domain, the optimum level is strictly 0 dBFS, although we often allow a safety margin of 3 dB (i.e., peak at –3 dBFS). The peak-hold feature, which is available on most digital meters, is extremely handy in these situations. It enables us to determine the exact amount of gain change straight after the first listen. When peak-hold is not available, we might have to listen to the material a few times while adjusting the level to optimum.

It does not end here. There are different ways to alter the level of the to-be-bounced material, and some are better than others. Here are the possible ways on an analog desk:

- **VCA group**—when available, this provides the quickest way to alter the levels of the original channels. Since VCA grouping does not involve an additional signal path, correct gain structure is maintained.
- **Channel faders**—a correct gain structure will also be maintained if all the channel faders of the original tracks are brought up. If the faders are motorized, there is usually a grouping or linking function that enables simultaneous movement of all the faders by the same number of decibels. If faders are not motorized, they should be adjusted one at a time, by the same number of decibels—a slow and imprecise method.
- **Group level**—since all the channels are sent to a group bus anyway, it is possible to bring down (and on many consoles up) the overall group level. While doing so is quick and easy, it does not comply with the correct gain structure concept, and usually results in additional noise (although sometimes this is unnoticeable).

In the realm of software mixers, the bouncing process is far more forgiving, as digital audio does not involve noisy components. The ways to alter the bounced material level are:

- **Master fader**—since most master faders in today's software mixers are scaling faders, they provide the quickest and the easiest way to alter the gain of the bounced signal. There is no quality penalty in using them.
- **Subgroup level**—routing all the original tracks to a bus and then altering the bus level. This is a longer method of achieving the same result, but valid just the same.
- **Fader grouping**—this involves grouping a set of track faders so they all move simultaneously. Here as well, there are no gain structure issues.

There is even more to consider. When we alter the level of the original tracks (as in the case of VCA grouping, channel faders, or fader grouping), we should take into account **post-fader aux sends**. If a snare is sent to a reverb emulator post-fader, there is no need to bring down the level of the reverb since it will drop down with respect to the snare fader. This is not the case when the send is pre-fader, where bringing down the snare level will not alter the reverb level, and so the reverb fader should be brought down as well. To summarize: setting the bounced material level should be done using the master fader on a software mixer, whereas, on a console, with a bounced mix that uses post-fader sends, it might be worth trading the addition of noise for simplicity and time, and using group levels.

Using the same principal of bouncing at mix levels, when individual stereo tracks are bounced they often suffer from **limited or shifted stereo width**, as they are bounced with the pan-pots positioned anywhere but the extremes. In the case of bounced synthesizer tracks, the same limited stereo width can be the result of a pan-spread control not set to full (like a unison pan-spread set to 50 percent). Stereo tracks that are not bounced panned to the extremes limit the mixing engineer's ability to control stereo aspects.

Setting the time range prior to bouncing often also requires some attention. It is worth remembering that most time-based effects will require additional time to fade away—if bouncing is limited to the original material's time range, reverbs and delays might cut abruptly. In addition, plugin delay compensation issues can lead to a bounced track that is out of sync with the original material. Figure A.3 shows an original snare hit (top) that was bounced with a reverb (bottom). There are three problems with the bounced version: it is lower in level than the original track, it is out of sync, and the reverb cuts at the end.
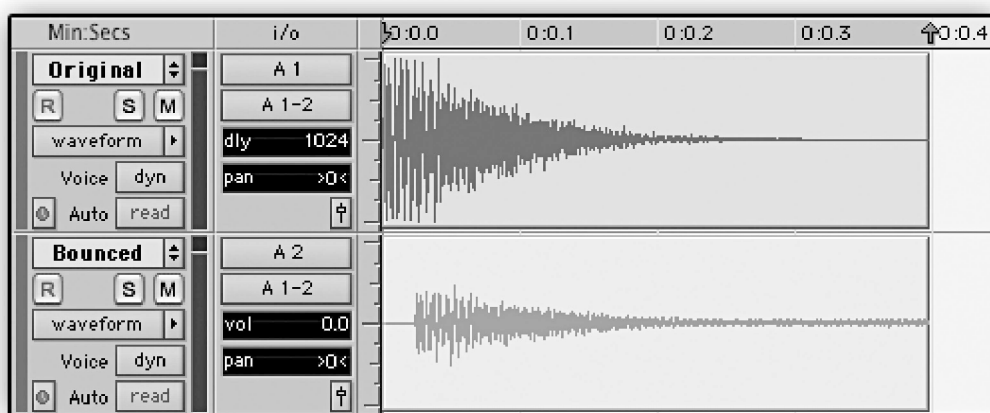
Figure A.3 Three problems caused by bouncing.

Once bouncing is complete, it is difficult to alter aspects of the original raw tracks in the bounced track; for example, the level balance between the kick and the snare. Yet, in most cases, the original tracks can be retrieved, mixed, and bounced again. Nearly every time we bounce, there is some quality penalty—noise in the analog domain,

digital distortion, dither noise, or dynamic range reduction (due to safety margins) can all be introduced during the process. While these unwanted additions might not be audible after the first instance of bouncing, they accumulate and become more profound when material is bounced time and again. Ideally, bouncing should not take place at all.

To summarize this section, Figure A.4 illustrates the recommended steps in the process of bouncing.
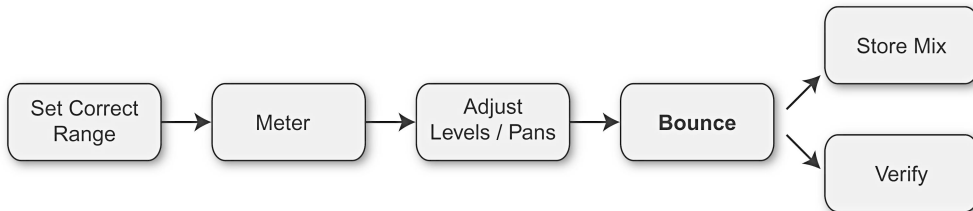
Figure A.4 Recommended steps in the process of bouncing.

APPENDIX B

Notes-to-frequencies chart

Octave		Note	Frequency		Octave		Note	Frequency
		C	16.35				C	523.25
		C#	17.32				C#	554.37
		D	18.35				D	587.33
		D#	19.45				D#	622.25
		E	20.6				E	659.26
		F	21.83				F	698.46
0		F#	23.12		5		F#	739.99
		G	24.5				G	783.99
		G#	25.96				G#	830.61
		A	27.5				A	880.00
		A#	29.14				A#	932.33
		B	30.87				B	987.77
		C	32.7				C	1046.5
		C#	34.65				C#	1108.73
		D	36.71				D	1174.66
		D#	38.89				D#	1244.51
		E	41.2 (Lowest E on standard bass guitar)				E	1318.51
		F	43.65				F	1396.91
1		F#	46.25		6		F#	1479.98
		G	49.00				G	1567.98
		G#	51.91				G#	1661.22
		A	55.00				A	1760.00
		A#	58.27				A#	1864.66
		B	61.74				B	1975.53
		C	65.41				C	2093.00
		C#	69.3				C#	2217.46
		D	73.42				D	2349.32
		D#	77.78				D#	2489.02
		E	82.41 (Lowest E on standard guitar)				E	2637.02
		F	87.31				F	2793.83
2		F#	92.5		7		F#	2959.96
		G	98.00				G	3135.96
		G#	103.83				G#	3322.44
		A	110.00				A	3520.00
		A#	116.54				A#	3729.31
		B	123.47				B	3951.07
		C	130.81				C	4186.01
		C#	138.59				C#	4434.92
		D	146.83				D	4698.64
		D#	155.56				D#	4978.03
		E	164.81				E	5274.04
		F	174.61				F	5587.65
3		F#	185.00		8		F#	5919.91
		G	196.00				G	6271.93
		G#	207.65				G#	6644.88
		A	220.00				A	7040.00
		A#	233.08				A#	7458.62
		B	246.94				B	7902.13
		C	261.63 (Middle C)				C	8372.02
		C#	277.18				C#	8869.84
		D	293.66				D	9397.27
		D#	311.13				D#	9956.06
		E	329.63				E	10548.08
		F	349.23				F	11175.3
4		F#	369.99		9		F#	11839.82
		G	392.00				G	12543.85
		G#	415.3				G#	13289.75
		A	**440.00** (Reference Pitch)				A	14080.00
		A#	466.16				A#	14917.24
		B	493.88				B	15804.27

APPENDIX C
Delay time chart

Delay times can be easily calculated. For time signatures involving quarter-note beats (4/4, 3/4, etc.), a quarter-note delay in seconds is calculated using 60/bpm. For example, at 120 bpm, a quarter-note delay would be 60/120, or 0.5 seconds, which is 500 ms. Milliseconds are far more common in delay units: using the calculation 60,000/bpm would give the result in milliseconds. A half-note delay is twice as long as a quarter-note delay for the same bpm, and an eighth-note delay is half the quarter-note delay. For example, a half-note delay at 120 bpm would be 1,000 ms, and an eighth-note delay would be 250 ms.

Simple note values

The following table includes delay times for various simple note values at 80–190 bpm. For faster bpm figures, simply halve the delay time of half the destination bpm. For example, a quarter-note delay at 240 bpm is 250 ms (half of 500 ms). For slower bpm figures, simply double the delay time of twice the destination bpm. For example, a quarter-note delay in 60 bpm is 1,000 ms (twice 500 ms). The delay time of a whole note in common time (4.4) is also the length of each bar.

bpm/note value	1/1	1/2	1/4	1/8	1/16
80	3,000.00	1,500.00	750.00	375.00	187.50
81	2,962.96	1,481.48	740.74	370.37	185.19
82	2,926.83	1,463.41	731.71	365.85	182.93
83	2,891.57	1,445.78	722.89	361.45	180.72
84	2,857.14	1,428.57	714.29	357.14	178.57
85	2,823.53	1,411.76	705.88	352.94	176.47
86	2,790.70	1,395.35	697.67	348.84	174.42
87	2,758.62	1,379.31	689.66	344.83	172.41
88	2,727.27	1,363.64	681.82	340.91	170.45
89	2,696.63	1,348.31	674.16	337.08	168.54
90	2,666.67	1,333.33	666.67	333.33	166.67
91	2,637.36	1,318.68	659.34	329.67	164.84
92	2,608.70	1,304.35	652.17	326.09	163.04
93	2,580.65	1,290.32	645.16	322.58	161.29

bpm/note value	1/1	1/2	1/4	1/8	1/16
94	2,553.19	1,276.60	638.30	319.15	159.57
95	2,526.32	1,263.16	631.58	315.79	157.89
96	2,500.00	1,250.00	625.00	312.50	156.25
97	2,474.23	1,237.11	618.56	309.28	154.64
98	2,448.98	1,224.49	612.24	306.12	153.06
99	2,424.24	1,212.12	606.06	303.03	151.52
100	2,400.00	1,200.00	600.00	300.00	150.00
101	2,376.24	1,188.12	594.06	297.03	148.51
102	2,352.94	1,176.47	588.24	294.12	147.06
103	2,330.10	1,165.05	582.52	291.26	145.63
104	2,307.69	1,153.85	576.92	288.46	144.23
105	2,285.71	1,142.86	571.43	285.71	142.86
106	2,264.15	1,132.08	566.04	283.02	141.51
107	2,242.99	1,121.50	560.75	280.37	140.19
108	2,222.22	1,111.11	555.56	277.78	138.89
109	2,201.83	1,100.92	550.46	275.23	137.61
110	2,181.82	1,090.91	545.45	272.73	136.36
111	2,162.16	1,081.08	540.54	270.27	135.14
112	2,142.86	1,071.43	535.71	267.86	133.93
113	2,123.89	1,061.95	530.97	265.49	132.74
114	2,105.26	1,052.63	526.32	263.16	131.58
115	2,086.96	1,043.48	521.74	260.87	130.43
116	2,068.97	1,034.48	517.24	258.62	129.31
117	2,051.28	1,025.64	512.82	256.41	128.21
118	2,033.90	1,016.95	508.47	254.24	127.12
119	2,016.81	1,008.40	504.20	252.10	126.05
120	2,000.00	1,000.00	500.00	250.00	125.00
121	1,983.47	991.74	495.87	247.93	123.97
122	1,967.21	983.61	491.80	245.90	122.95
123	1,951.22	975.61	487.80	243.90	121.95
124	1,935.48	967.74	483.87	241.94	120.97
125	1,920.00	960.00	480.00	240.00	120.00
126	1,904.76	952.38	476.19	238.10	119.05
127	1,889.76	944.88	472.44	236.22	118.11
128	1,875.00	937.50	468.75	234.38	117.19
129	1,860.47	930.23	465.12	232.56	116.28
130	1,846.15	923.08	461.54	230.77	115.38
131	1,832.06	916.03	458.02	229.01	114.50
132	1,818.18	909.09	454.55	227.27	113.64
133	1,804.51	902.26	451.13	225.56	112.78
134	1,791.04	895.52	447.76	223.88	111.94
135	1,777.78	888.89	444.44	222.22	111.11
136	1,764.71	882.35	441.18	220.59	110.29
137	1,751.82	875.91	437.96	218.98	109.49
138	1,739.13	869.57	434.78	217.39	108.70
139	1,726.62	863.31	431.65	215.83	107.91
140	1,714.29	857.14	428.57	214.29	107.14
141	1,702.13	851.06	425.53	212.77	106.38
142	1,690.14	845.07	422.54	211.27	105.63
143	1,678.32	839.16	419.58	209.79	104.90

bpm/note value	1/1	1/2	1/4	1/8	1/16
144	1,666.67	833.33	416.67	208.33	104.17
145	1,655.17	827.59	413.79	206.90	103.45
146	1,643.84	821.92	410.96	205.48	102.74
147	1,632.65	816.33	408.16	204.08	102.04
148	1,621.62	810.81	405.41	202.70	101.35
149	1,610.74	805.37	402.68	201.34	100.67
150	1,600.00	800.00	400.00	200.00	100.00
151	1,589.40	794.70	397.35	198.68	99.34
152	1,578.95	789.47	394.74	197.37	98.68
153	1,568.63	784.31	392.16	196.08	98.04
154	1,558.44	779.22	389.61	194.81	97.40
155	1,548.39	774.19	387.10	193.55	96.77
156	1,538.46	769.23	384.62	192.31	96.15
157	1,528.66	764.33	382.17	191.08	95.54
158	1,518.99	759.49	379.75	189.87	94.94
159	1,509.43	754.72	377.36	188.68	94.34
160	1,500.00	750.00	375.00	187.50	93.75
161	1,490.68	745.34	372.67	186.34	93.17
162	1,481.48	740.74	370.37	185.19	92.59
163	1,472.39	736.20	368.10	184.05	92.02
164	1,463.41	731.71	365.85	182.93	91.46
165	1,454.55	727.27	363.64	181.82	90.91
166	1,445.78	722.89	361.45	180.72	90.36
167	1,437.13	718.56	359.28	179.64	89.82
168	1,428.57	714.29	357.14	178.57	89.29
169	1,420.12	710.06	355.03	177.51	88.76
170	1,411.76	705.88	352.94	176.47	88.24
171	1,403.51	701.75	350.88	175.44	87.72
172	1,395.35	697.67	348.84	174.42	87.21
173	1,387.28	693.64	346.82	173.41	86.71
174	1,379.31	689.66	344.83	172.41	86.21
175	1,371.43	685.71	342.86	171.43	85.71
176	1,363.64	681.82	340.91	170.45	85.23
177	1,355.93	677.97	338.98	169.49	84.75
178	1,348.31	674.16	337.08	168.54	84.27
179	1,340.78	670.39	335.20	167.60	83.80
180	1,333.33	666.67	333.33	166.67	83.33
181	1,325.97	662.98	331.49	165.75	82.87
182	1,318.68	659.34	329.67	164.84	82.42
183	1,311.48	655.74	327.87	163.93	81.97
184	1,304.35	652.17	326.09	163.04	81.52
185	1,297.30	648.65	324.32	162.16	81.08
186	1,290.32	645.16	322.58	161.29	80.65
187	1,283.42	641.71	320.86	160.43	80.21
188	1,276.60	638.30	319.15	159.57	79.79
189	1,269.84	634.92	317.46	158.73	79.37
190	1,263.16	631.58	315.79	157.89	78.95

Triplets

To calculate the delay time of a triplet note, simply multiply the delay time of the simple note value by 2/3 (0.6667). The following chart displays the delay times for triplet delays.

bpm/note value	1/1 T	1/2 T	1/4 T	1/8 T	1/16 T
80	2,000.00	1,000.00	500.00	250.00	125.00
81	1,975.31	987.65	493.83	246.91	123.46
82	1,951.22	975.61	487.80	243.90	121.95
83	1,927.71	963.86	481.93	240.96	120.48
84	1,904.76	952.38	476.19	238.10	119.05
85	1,882.35	941.18	470.59	235.29	117.65
86	1,860.47	930.23	465.12	232.56	116.28
87	1,839.08	919.54	459.77	229.89	114.94
88	1,818.18	909.09	454.55	227.27	113.64
89	1,797.75	898.88	449.44	224.72	112.36
90	1,777.78	888.89	444.44	222.22	111.11
91	1,758.24	879.12	439.56	219.78	109.89
92	1,739.13	869.57	434.78	217.39	108.70
93	1,720.43	860.22	430.11	215.05	107.53
94	1,702.13	851.06	425.53	212.77	106.38
95	1,684.21	842.11	421.05	210.53	105.26
96	1,666.67	833.33	416.67	208.33	104.17
97	1,649.48	824.74	412.37	206.19	103.09
98	1,632.65	816.33	408.16	204.08	102.04
99	1,616.16	808.08	404.04	202.02	101.01
100	1,600.00	800.00	400.00	200.00	100.00
101	1,584.16	792.08	396.04	198.02	99.01
102	1,568.63	784.31	392.16	196.08	98.04
103	1,553.40	776.70	388.35	194.17	97.09
104	1,538.46	769.23	384.62	192.31	96.15
105	1,523.81	761.90	380.95	190.48	95.24
106	1,509.43	754.72	377.36	188.68	94.34
107	1,495.33	747.66	373.83	186.92	93.46
108	1,481.48	740.74	370.37	185.19	92.59
109	1,467.89	733.94	366.97	183.49	91.74
110	1,454.55	727.27	363.64	181.82	90.91
111	1,441.44	720.72	360.36	180.18	90.09
112	1,428.57	714.29	357.14	178.57	89.29
113	1,415.93	707.96	353.98	176.99	88.50
114	1,403.51	701.75	350.88	175.44	87.72
115	1,391.30	695.65	347.83	173.91	86.96
116	1,379.31	689.66	344.83	172.41	86.21
117	1,367.52	683.76	341.88	170.94	85.47
118	1,355.93	677.97	338.98	169.49	84.75
119	1,344.54	672.27	336.13	168.07	84.03
120	1,333.33	666.67	333.33	166.67	83.33
121	1,322.31	661.16	330.58	165.29	82.64
122	1,311.48	655.74	327.87	163.93	81.97
123	1,300.81	650.41	325.20	162.60	81.30

bpm/note value	1/1 T	1/2 T	1/4 T	1/8 T	1/16 T
124	1,290.32	645.16	322.58	161.29	80.65
125	1,280.00	640.00	320.00	160.00	80.00
126	1,269.84	634.92	317.46	158.73	79.37
127	1,259.84	629.92	314.96	157.48	78.74
128	1,250.00	625.00	312.50	156.25	78.13
129	1,240.31	620.16	310.08	155.04	77.52
130	1,230.77	615.38	307.69	153.85	76.92
131	1,221.37	610.69	305.34	152.67	76.34
132	1,212.12	606.06	303.03	151.52	75.76
133	1,203.01	601.50	300.75	150.38	75.19
134	1,194.03	597.01	298.51	149.25	74.63
135	1,185.19	592.59	296.30	148.15	74.07
136	1,176.47	588.24	294.12	147.06	73.53
137	1,167.88	583.94	291.97	145.99	72.99
138	1,159.42	579.71	289.86	144.93	72.46
139	1,151.08	575.54	287.77	143.88	71.94
140	1,142.86	571.43	285.71	142.86	71.43
141	1,134.75	567.38	283.69	141.84	70.92
142	1,126.76	563.38	281.69	140.85	70.42
143	1,118.88	559.44	279.72	139.86	69.93
144	1,111.11	555.56	277.78	138.89	69.44
145	1,103.45	551.72	275.86	137.93	68.97
146	1,095.89	547.95	273.97	136.99	68.49
147	1,088.44	544.22	272.11	136.05	68.03
148	1,081.08	540.54	270.27	135.14	67.57
149	1,073.83	536.91	268.46	134.23	67.11
150	1,066.67	533.33	266.67	133.33	66.67
151	1,059.60	529.80	264.90	132.45	66.23
152	1,052.63	526.32	263.16	131.58	65.79
153	1,045.75	522.88	261.44	130.72	65.36
154	1,038.96	519.48	259.74	129.87	64.94
155	1,032.26	516.13	258.06	129.03	64.52
156	1,025.64	512.82	256.41	128.21	64.10
157	1,019.11	509.55	254.78	127.39	63.69
158	1,012.66	506.33	253.16	126.58	63.29
159	1,006.29	503.14	251.57	125.79	62.89
160	1,000.00	500.00	250.00	125.00	62.50
161	993.79	496.89	248.45	124.22	62.11
162	987.65	493.83	246.91	123.46	61.73
163	981.60	490.80	245.40	122.70	61.35
164	975.61	487.80	243.90	121.95	60.98
165	969.70	484.85	242.42	121.21	60.61
166	963.86	481.93	240.96	120.48	60.24
167	958.08	479.04	239.52	119.76	59.88
168	952.38	476.19	238.10	119.05	59.52
169	946.75	473.37	236.69	118.34	59.17
170	941.18	470.59	235.29	117.65	58.82
171	935.67	467.84	233.92	116.96	58.48
172	930.23	465.12	232.56	116.28	58.14
173	924.86	462.43	231.21	115.61	57.80

bpm/note value	1/1 T	1/2 T	1/4 T	1/8 T	1/16 T
174	919.54	459.77	229.89	114.94	57.47
175	914.29	457.14	228.57	114.29	57.14
176	909.09	454.55	227.27	113.64	56.82
177	903.95	451.98	225.99	112.99	56.50
178	898.88	449.44	224.72	112.36	56.18
179	893.85	446.93	223.46	111.73	55.87
180	888.89	444.44	222.22	111.11	55.56
181	883.98	441.99	220.99	110.50	55.25
182	879.12	439.56	219.78	109.89	54.95
183	874.32	437.16	218.58	109.29	54.64
184	869.57	434.78	217.39	108.70	54.35
185	864.86	432.43	216.22	108.11	54.05
186	860.22	430.11	215.05	107.53	53.76
187	855.61	427.81	213.90	106.95	53.48
188	851.06	425.53	212.77	106.38	53.19
189	846.56	423.28	211.64	105.82	52.91
190	842.11	421.05	210.53	105.26	52.63

Dotted notes

A dotted note duration is calculated by adding half of the note duration to the full note duration. So to calculate the delay time of a dotted note, all we need to do is multiply the delay time of the simple note value by 1.5. The following chart displays the delay time for dotted note delays.

bpm/note value	1/1D	1/2D	1/4D	1/8D	1/16D
80	4,500.00	2,250.00	1,125.00	562.50	281.25
81	4,444.44	2,222.22	1,111.11	555.56	277.78
82	4,390.24	2,195.12	1,097.56	548.78	274.39
83	4,337.35	2,168.67	1,084.34	542.17	271.08
84	4,285.71	2,142.86	1,071.43	535.71	267.86
85	4,235.29	2,117.65	1,058.82	529.41	264.71
86	4,186.05	2,093.02	1,046.51	523.26	261.63
87	4,137.93	2,068.97	1,034.48	517.24	258.62
88	4,090.91	2,045.45	1,022.73	511.36	255.68
89	4,044.94	2,022.47	1,011.24	505.62	252.81
90	4,000.00	2,000.00	1,000.00	500.00	250.00
91	3,956.04	1,978.02	989.01	494.51	247.25
92	3,913.04	1,956.52	978.26	489.13	244.57
93	3,870.97	1,935.48	967.74	483.87	241.94
94	3,829.79	1,914.89	957.45	478.72	239.36
95	3,789.47	1,894.74	947.37	473.68	236.84
96	3,750.00	1,875.00	937.50	468.75	234.38
97	3,711.34	1,855.67	927.84	463.92	231.96
98	3,673.47	1,836.73	918.37	459.18	229.59
99	3,636.36	1,818.18	909.09	454.55	227.27
100	3,600.00	1,800.00	900.00	450.00	225.00

bpm/note value	1/1D	1/2D	1/4D	1/8D	1/16D
101	3,564.36	1,782.18	891.09	445.54	222.77
102	3,529.41	1,764.71	882.35	441.18	220.59
103	3,495.15	1,747.57	873.79	436.89	218.45
104	3,461.54	1,730.77	865.38	432.69	216.35
105	3,428.57	1,714.29	857.14	428.57	214.29
106	3,396.23	1,698.11	849.06	424.53	212.26
107	3,364.49	1,682.24	841.12	420.56	210.28
108	3,333.33	1,666.67	833.33	416.67	208.33
109	3,302.75	1,651.38	825.69	412.84	206.42
110	3,272.73	1,636.36	818.18	409.09	204.55
111	3,243.24	1,621.62	810.81	405.41	202.70
112	3,214.29	1,607.14	803.57	401.79	200.89
113	3,185.84	1,592.92	796.46	398.23	199.12
114	3,157.89	1,578.95	789.47	394.74	197.37
115	3,130.43	1,565.22	782.61	391.30	195.65
116	3,103.45	1,551.72	775.86	387.93	193.97
117	3,076.92	1,538.46	769.23	384.62	192.31
118	3,050.85	1,525.42	762.71	381.36	190.68
119	3,025.21	1,512.61	756.30	378.15	189.08
120	3,000.00	1,500.00	750.00	375.00	187.50
121	2,975.21	1,487.60	743.80	371.90	185.95
122	2,950.82	1,475.41	737.70	368.85	184.43
123	2,926.83	1,463.41	731.71	365.85	182.93
124	2,903.23	1,451.61	725.81	362.90	181.45
125	2,880.00	1,440.00	720.00	360.00	180.00
126	2,857.14	1,428.57	714.29	357.14	178.57
127	2,834.65	1,417.32	708.66	354.33	177.17
128	2,812.50	1,406.25	703.13	351.56	175.78
129	2,790.70	1,395.35	697.67	348.84	174.42
130	2,769.23	1,384.62	692.31	346.15	173.08
131	2,748.09	1,374.05	687.02	343.51	171.76
132	2,727.27	1,363.64	681.82	340.91	170.45
133	2,706.77	1,353.38	676.69	338.35	169.17
134	2,686.57	1,343.28	671.64	335.82	167.91
135	2,666.67	1,333.33	666.67	333.33	166.67
136	2,647.06	1,323.53	661.76	330.88	165.44
137	2,627.74	1,313.87	656.93	328.47	164.23
138	2,608.70	1,304.35	652.17	326.09	163.04
139	2,589.93	1,294.96	647.48	323.74	161.87
140	2,571.43	1,285.71	642.86	321.43	160.71
141	2,553.19	1,276.60	638.30	319.15	159.57
142	2,535.21	1,267.61	633.80	316.90	158.45
143	2,517.48	1,258.74	6,29.37	314.69	157.34
144	2,500.00	1,250.00	625.00	312.50	156.25
145	2,482.76	1,241.38	620.69	310.34	155.17
146	2,465.75	1,232.88	616.44	308.22	154.11
147	2,448.98	1,224.49	612.24	306.12	153.06
148	2,432.43	1,216.22	608.11	304.05	152.03
149	2,416.11	1,208.05	604.03	302.01	151.01
150	2,400.00	1,200.00	600.00	300.00	150.00

bpm/note value	1/1D	1/2D	1/4D	1/8D	1/16D
151	2,384.11	1,192.05	596.03	298.01	149.01
152	2,368.42	1,184.21	592.11	296.05	148.03
153	2,352.94	1,176.47	588.24	294.12	147.06
154	2,337.66	1,168.83	584.42	292.21	146.10
155	2,322.58	1,161.29	580.65	290.32	145.16
156	2,307.69	1,153.85	576.92	288.46	144.23
157	2,292.99	1,146.50	573.25	286.62	143.31
158	2,278.48	1,139.24	569.62	284.81	142.41
159	2,264.15	1,132.08	566.04	283.02	141.51
160	2,250.00	1,125.00	562.50	281.25	140.63
161	2,236.02	1,118.01	559.01	279.50	139.75
162	2,222.22	1,111.11	555.56	277.78	138.89
163	2,208.59	1,104.29	552.15	276.07	138.04
164	2,195.12	1,097.56	548.78	274.39	137.20
165	2,181.82	1,090.91	545.45	272.73	136.36
166	2,168.67	1,084.34	542.17	271.08	135.54
167	2,155.69	1,077.84	538.92	269.46	134.73
168	2,142.86	1,071.43	535.71	267.86	133.93
169	2,130.18	1065.09	532.54	266.27	133.14
170	2,117.65	1,058.82	529.41	264.71	132.35
171	2,105.26	1,052.63	526.32	263.16	131.58
172	2,093.02	1,046.51	523.26	261.63	130.81
173	2,080.92	1,040.46	520.23	260.12	130.06
174	2,068.97	1,034.48	517.24	258.62	129.31
175	2,057.14	1,028.57	514.29	257.14	128.57
176	2,045.45	1,022.73	511.36	255.68	127.84
177	2,033.90	1,016.95	508.47	254.24	127.12
178	2,022.47	1,011.24	505.62	252.81	126.40
179	2,011.17	1,005.59	502.79	251.40	125.70
180	2,000.00	1,000.00	500.00	250.00	125.00
181	1,988.95	994.48	497.24	248.62	124.31
182	1,978.02	989.01	494.51	247.25	123.63
183	1,967.21	983.61	491.80	245.90	122.95
184	1,956.52	978.26	489.13	244.57	122.28
185	1,945.95	972.97	486.49	243.24	121.62
186	1,935.48	967.74	483.87	241.94	120.97
187	1,925.13	962.57	481.28	240.64	120.32
188	1,914.89	957.45	478.72	239.36	119.68
189	1,904.76	952.38	476.19	238.10	119.05
190	1,894.74	947.37	473.68	236.84	118.42

Index

Note: Locators in *italics* represent figures in the text.